When Do Sin Taxes Increase Welfare?*

Tuomas Kosonen†, Sami Jysmä‡, and Riikka Savolainen§

March 21, 2026

Abstract

This paper studies the conditions under which sin taxes affect consumption and welfare. We organize the paper around a theoretical framework that characterizes how consumers' substitution preferences between two goods relate to the elasticity of consumption. The model yields a highly convex pattern in which demand elasticities are very small in general, but large when the two goods are very close substitutes. The theoretical framework also presents a welfare analysis of sin taxes, showing the conditions under which taxes increase welfare the most. These conditions are related to the magnitude of elasticities and externalities. Empirically, we analyze a Finnish sin tax scheme on sweets, soda, and ice cream that provides quasi-experimental variation through multiple reforms. We use product- and store-level data on sales and prices containing hundreds of millions of observations. We also provide survey evidence on substitution preferences across categories of goods. Our estimated consumption elasticities align well with the theoretical framework: consumption elasticity is high only for sugary soft drinks and juices, for which the non-taxed close substitute is sugar-free soft drinks/juices. Finally, we provide a meta-analysis of the literature on consumption elasticities and show that the elasticity estimates on average align with our theoretical framework.

Key words: excise taxes, sin tax, consumption, substitution, sweets, soda

JEL Codes: H2, I18

*We would like to thank Salvatore Barbaro, Raj Chetty, Stefano DellaVigna, John Friedman, David Frisvold, Nathan Hendren, Jim Poterba, Emmanuel Saez, Dmitry Taubinsky, Otto Toivanen and Danny Yagan as well as participants at numerous conferences and seminars for helpful comments. Miska Raivio and Katri Väisänen provided very helpful research assistance. We are grateful for the S-group for providing us their data and the Research Council of Finland for funding (grant no. 346253 and 328177). Sami Jysmä acknowledges support from Palkansaajasäätiö.

† VATT Institute for Economic Research, Finnish Centre of Excellence in Tax Systems Research FIT, tuomas.kosonen@vatt.fi.

‡ Labour Institute for Economic Research, Tampere University and FIT, sami.jysma@labore.fi

§ Swansea University, r.h.m.savolainen@swansea.ac.uk.

1 Introduction

Sin taxes that target goods containing sugar, alcohol or nicotine are seen as important tools to reduce socially harmful consumption (Andreyeva et al. (2010), Cawley et al. (2019) and Allcott et al. (2019b)). Among other taxes that target specific consumption goods are reduced VAT rates in the European Union, CO_2 dependent vehicle taxation and high excise duties for specific goods. A central question is, under what conditions these targeted taxes change the consumption of goods they target and hence affect welfare (see e.g. Allcott et al. (2019a)), a question that is central to the design of effective sin taxes, and other targeted consumption taxes. The empirical literature on sin taxes is not very helpful on this as it provides an extremely wide range of consumption elasticities (see e.g. Andreyeva et al. (2010), Cawley et al. (2019), Allcott et al. (2019a)). In particular, it does not provide an analytical answer as to why the elasticity range is so wide preventing policymakers from knowing what to expect when setting new sin taxes or other targeted consumption taxes. Moreover, as a consequence the conditions under which sin taxes increase welfare the most are not well understood.

In this paper we revisit these questions by analyzing a sin tax scheme with multiple reforms in a comprehensive scanner data containing around 500 million observations on the sales and prices of sin goods and other similar goods. We also provide a meta-analysis on papers providing reduced form analysis of targeted consumption taxes to see how our findings are generalizable to the broader literature.

We organize our analysis through a stylized framework that provides analytical predictions on the conditions when demand elasticity is high and when it is close to zero. The central idea is to consider a choice situation between a few goods that have varying levels of substitutability between them. The government then sets consumption taxes to selected goods and consumers decide to what extent they consume those goods or other similar and non-taxed goods. The basic framework has two goods and a representative consumer, which can be extended to multiple goods. The two goods share two-dimensional characteristics, which is a structure in an otherwise standard model allowing us to state within the model how similar, or substitutable, the two goods are. We

then investigate how the substitutability of goods relate to the demand elasticity of the taxed goods.

The model predicts an intriguing pattern between the substitutability of goods and demand elasticity. This relationship is highly non-convex such that when the two goods are fairly substitutable the demand elasticity stays close to zero, and only when the two goods are highly substitutable the demand elasticity is high, increasing exponentially with the substitutability index. Thus, the framework predicts analytically when, based on substitutability, targeted consumption taxes, such as sin taxes, strongly shift consumption and when they do not.

We also extend the theoretical framework to analyze the welfare implications of taxing goods whose consumption generates externalities.[1] We focus on the question, under what conditions taxes on sin goods increase welfare the most. To achieve this we present the Marginal Value of Public Funds, or MVPF formulas for sin taxes and numerical illustrations for it. For completeness, we derive also optimal commodity tax result. We show how the size of externalities, the demand elasticity and existing tax level are related to the MVPF. Welfare from sin taxes increases when substitutability between goods, and thus demand elasticity, is high and the substitute good creates less externalities or internalities, that is, in the case of sin taxes the good is healthier.

We then study empirically a sin tax scheme in Finland that has been reformed several times. This gives us rich variation in the level of sin taxes, and in how close substitutes are left outside of the reforms. The sin tax scheme was introduced in the beginning of January 2011 as a weight based tax on solid products (0.75€/kg) or volume based tax on drinks (0.075€/l). It applied to products in product categories containing sweets, soft drinks or ice creams and did not apply at all to other products that nevertheless contain added sugar or chocolate, such as cookies, sweet pastries or yoghurts. Thus, the sin tax scheme did not apply to products that can be seen as substitutes to the taxed products. The tax level was increased in 2012, and in 2017 the tax was removed for sweets and ice creams. In 2014, the tax rates of soft drinks were separated such that the tax rate

[1]The welfare analysis is in the spirit of that originally presented by Ramsey (1927) and Pigou (1920).

of sugary liquids (0.22€/l) became twice that of sugar-free liquids (0.11€/l). The 2014 reform treats very substitutable products, such as sugary and non-sugary soda from the same producer, differently. In summary, different reforms in the sin tax scheme leave potential substitution goods that have varying levels of substitutability with the taxed goods untaxed.

We have access to extremely large product-level data for the taxed and non-taxed goods. The unit of observation is a product sold in a store in a week. At this level, we observe how many products were sold and at what price. In total, we have about 500 million observations for tens of thousands of products. The data originate from S-group, which is the largest grocery store chain in Finland having close to 50% market share. Hence, although the data are sampled non-randomly, they describe in a meaningful way the consumption of goods targeted by the tax scheme in Finland. The data allow us to provide precise estimates and also investigate the role of potentially confounding factors, such as, overall trends in grocery consumption or the seasonality of the taxed products.

We also develop a survey that measures substitution preferences between goods allowing us to determine what the closest non-taxed substitute categories for taxed goods are in each reform. It is an online survey where we ask questions about pairs of products that are randomly selected from a large pool of product-pairs. The main question of interest asks how close substitutes a pair of products are on a scale from 0 to 10 with a higher number indicating closer substitutes.[2] The survey yields a substitution index between product groups. This allows us to determine which product groups are the closest potential substitution groups to the taxed products, and which are more distant goods that can be used as control products.

To estimate the impact of different reforms on consumer prices and quantities of goods sold, we provide graphical examination of trends and event-studies for the taxed goods, their potential substitution goods and control goods. We select the potential substitution and control groups by taking into account how close substitutes they are to the taxed products based on our survey results. We study substitution by examining to what extent

[2]We also build other measures for substitutability that utilize A) EU tariff goods tree or B) product or store information, e.g. nutrients or distance between products in a typical store.

the sales of the potential substitution groups respond to tax changes of the taxed goods relative to their control group. We observe high seasonality in the sales of products.

Our results for tax inclusive consumer prices show that for tax increases the sweets tax was either fully or even slightly over-shifted to prices. In total, the tax increases in 2011 and 2012 increased the prices of sweets by about 10%, and so did the 2014 reform for sugary drinks. In contrast, abolishing the tax in 2017 for sweets and ice cream led to a less than full pass-through with about 10% reduction in prices on average resulting in slight asymmetric tax incidence (see Benzarti et al. (2020)). An important observation is that prices change on average by roughly similar amounts relative to a control group in different reforms.

The results for quantities sold show two very distinct results. First, we estimate quite precise zero change in quantities of sweets, ice cream or soda using tax increases that treated similar products the same, for example sugary and sugar-free soda. The same applies to abolishing the tax in 2017 for sweets and ice cream. Second, the tax increase for sugary drinks in 2014 that left sugar-free drinks unaffected led to a drastic shift in quantities sold such that sales of sugary drinks fell at the reform while those of sugar-free drinks increased. When examining soft drinks as a whole, the reform in 2014 does not seem to have an effect on their total sales. Moreover, the estimate for sales of soft drinks from the 2011 reform that treated sugary and non-sugary drinks equally is essentially zero.

These estimates are then used to derive average price elasticities at the average levels of prices and demand. The tax changes for sweets, chocolate and ice-cream that left cookies as the closest substitute not facing tax changes yields elasticities below 0.1 in absolute value. On the other hand, for sugary soda that had as the closest substitute non-sugary soda and juices we estimate elasticities of over 0.5 in absolute value.

Our next step is to combine the theory framework and the quasi-experimental estimates. We align our estimated consumption elasticities for different product groups with the substitution index from our survey. We also predict with our model how large demand elasticities we would expect based on the same substitution index. We then compare the

model predictions and our estimated elasticities for different product pairs. The model predicts that, for sweets and ice-cream, which have as the closest non-taxed substitute cookies, the demand elasticity should be quite low. Instead, for sugary soda that has as the closest substitute non-sugary soda, the demand elasticity should be at a clearly higher level in absolute value. In summary, the pattern of predicted elasticities is highly convex, when ordered in an increasing order based on the substitutability to the non-taxed goods. This prediction from the model aligns well with the pattern of our empirical elasticity estimates. Thus, our theory framework seems to perform well in predicting the estimated demand elasticities based only on information about how substitutable different goods are.

Our second complimentary approach for studying empirically the link between consumption elasticities and substitutability to non-taxed goods is a meta-analysis to earlier literature. To align the estimates from earlier literature with the prediction from the model requires us to develop a substitution index. The substitution index we use is based on a product category tree from the TARIC (Integrated Tariff of the European Communities), which is available to virtually all commodities. The index is based on how many nods in the tree one needs to cross to get from one product to the potential substitute group, and therefore, how many goods are at a comparable level in the tariff tree. The less nods the higher is the substitutability. We complement the TARIC information with information about geographical proximity (in case of cross-border shopping) and other assessments of substitutes for fuel and services.

We include into the meta-analysis all articles that provide credible reduced form consumption elasticity estimates and for which we can assess the substitution index. The included studies have found a large range of consumption elasticities. We then group the estimates from literature into four categories based on the availability of substitutes that do not face tax changes; very low (VAT changes for broad goods or services), low (gasoline or electricity taxes), intermediate (alcohol taxes with cross-border shopping) and high (cigarettes vs e-cigarettes, sweetened vs non-sweetened beverages) substitutability. We use the average substitution index of the category to place them on an ordinal scale.

The result of the meta-analysis is that the average of estimated consumption elasticities in the literature are low when substitution index is low and that consumption elasticities increase in a highly convex shape with the substitution index. The result resembles remarkably well the pattern of elasticities from our sweets tax reforms, and aligns well with the prediction from our theory framework. Note that a selection of individual estimates, when not grouped into averages might not fit well with the theory model. Thus, the available of non-taxed substitute may not be the only factor explaining the wide range of elasticities in the literature.

Other reduced form estimates on consumption seem to align with our theory framework, although we did not include them in the meta-analysis due to the lack of estimated elasticities. For example, in the high substitutability category would be Harju et al. (2020) who study taxing new cars with CO2 based taxes. The substitutability in this case is high as one can substitute a higher CO2 emissions car with lower CO2 emissions car or a new car with used car, that is, substitute a higher taxed car with lower taxed car. The evidence points that the reform introducing CO2 emissions based car tax dramatically shifted the consumption away from high tax and CO2 emission cars towards low tax and CO2 emission cars. Johansson et al. (2014) study how a Finnish alcohol tax cut affected alcohol sales in neighboring Sweden as a function of distance to the Finnish border. They find large changes in alcohol sales in Sweden close to the Finnish border, where alcohol on the Finnish side is a close substitute to alcohol on the Swedish side, but these effects diminish rapidly further off the border.

Sin taxes have also been studied extensively using structural estimates (see e.g. Chan (2006), Dubois et al. (2020), Miravete et al. (2020)). Some of these structural papers estimate models that share a feature with our framework in that they include product characteristics as something consumers care about, for example Dubois et al. (2014) develop and estimate a theoretical framework, which nests a characteristics model from the family of Gorman-Lancaster models (Gorman (1980) and Lancaster (1966)). The advantage of structural models over our approach is that they allow to estimate the whole demand system, and therefore to also do counterfactual analysis of the effects of

different kinds of policy changes. Our estimates do not produce a demand curve, they are instead average estimates for the elasticities in the existing environment. We think that the advantage of our approach is to provide a credibly causal estimate with transparent event-study evidence supporting the assumptions. Moreover, our study complements the structural literature in providing analytical results and quasi-experimental support for the characteristics-based consumption models.

We contribute also to the literature on consumption tax incidence by estimating price incidences of excise taxes. From earlier studies Berardi et al. (2016) study the soda tax pass-through in France using a data set similar to ours, with outlet-level super market prices. They find similarly to us that the pass-through of the tax to prices was not immediate, and instead it took six months to reach the full pass-through. Also similar to our estimates for soft drinks, Grogger (2017) documents over-shifting into the prices of soda prices as the response to the introduction of a soda tax in Mexico. Similarly, Bergman and Hansen (2019) find that excise taxes are over-shifted to sodas when analyzing Danish micro-data.

One of our price incidence results is the asymmetric pass-through to prices from tax increases vs tax decreases, similar to Benzarti et al. (2020) for VAT. Bergman and Hansen (2019) also document that their pass-through rates are asymmetric so that excise tax increases are overshifted more than tax cuts. Moreover, we find smaller pass-through rates when the tax was larger relative to the pre-reform price. Agrawal and Hoyt (2019) show that over-shifting of taxes onto prices may also be a result of the presence of substitute goods.

The paper proceeds as follows: Section 2 presents our stylized theoretical model, Section 3 describes the institutions more precisely, Section 4 describes the data, Section 5 shows the empirical results for the sweets tax scheme and how they fit our theoretical framework, Section 6 discusses how empirical estimates from previous literature and the substitution index we develop align with our theoretical framework, while Section 7 concludes.

2 Conceptual framework

2.1 A model with characteristics

In this section, we set up a stylized framework to study the relationship between the similarity, or substitutability, of goods and the demand elasticity. With the framework we want to characterize a simple choice situation where consumer chooses from a limited set of goods to consume, a situation that could occur for example when one considers buying a sin good or a healthier alternative.

The framework is otherwise a standard model of consumption over a vector of goods, but to build a structure allowing us to investigate substitutability, we include characteristics of goods into the framework.[3] Consumers value the characteristics of goods in their consumption choices. The characteristics relate to the substitutability of goods so that the more similar characteristics the goods have, the more substitutable they are. Intuitively, the demand elasticity is a function of the extent to which consumers are willing to substitute one good for another, which is determined by the characteristics of the goods.

Below, we relate the characteristics-based measure of substitutability to empirical substitution indices. In Section 5.4 we calibrate the model to our empirical setting using substitution indices and see how well the empirical results align with the predictions from the model. The goods in the model are empirically assumed to come from different product categories, so the goods correspond to goods categories.

A basic version of our model is a two-good model where the goods have two characteristics and there are no income effects. Thus, the thought experiment is to consider a choice situation between, for example, a sin good and another good. The model produces a highly convex relationship between the demand elasticity of a good and the substitutability of the two goods: for low to relatively high levels of substitutability, the elasticity is small, but for highly substitutable goods, the elasticity is considerably higher. We use this relationship to organize our empirical work. The treatment of the model in the main text emphasizes the intuition behind our results, but more formal and detailed

[3]The characteristics space follows the spirit of Gorman (1980) and Lancaster (1966), but here we consider characteristics to be composite objects capturing subjective substitution preferences.

derivations together with a model with more than two goods and income effects can be found in Appendix B.

2.1.1 Setup of the model

The model is set up to consider the consumption choices of a single agent between two goods containing potentially different sets of characteristics. A utility-maximizing agent chooses the consumption of two goods $x = (x_1, x_2)'$ such that the budget constraint is satisfied. The agent has a strictly concave utility function $u(c)$ over bundles of characteristics $c = (c_1, c_2)'$, where characteristics provide a description of the consumption bundle. The characteristics are subjective composites that the consumer values, so they could include, for example, brand and packaging in one characteristic and taste in another. Good i has a characteristics vector $c^i = (c_{1i}, c_{2i})'$ which describes the characteristics of the good. The main idea is that a model involving characteristics of goods provides a way to measure similarity, and hence substitutability, without directly estimating cross-elasticities.

We assume, as is standard in this class of models, that consuming a bundle of goods x provides characteristics through a linear relationship:

$$c = Cx,$$

where $C = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{bmatrix}$ is a matrix of the characteristics of the goods. Put in another way, the amount of characteristics c the agent obtains from a given bundle of goods x, is given by

$$c = (c_{11}x_1 + c_{12}x_2, c_{21}x_1 + c_{22}x_2).$$

Hence the amount of a given characteristic in a bundle is a weighted sum over goods. This is natural for additive characteristics: if one characteristic is sugar content, the bundle's sugar is a sum over the sugar content of the two consumed goods.

To parameterize substitutability in this strict two-good setup, let good 1 have fixed

characteristics vector c^1, and let good 2 become more similar to good 1 as α increases:

$$c^2(\alpha) = \alpha c^1 + (1-\alpha)c^{2,0}, \qquad \alpha \in [0,1],$$

where $c^{2,0}$ is the baseline characteristics vector of good 2. Equivalently, the characteristics matrix can be written as $C(\alpha) = [c^1, c^2(\alpha)]$. Thus, when $\alpha = 0$, good 2 is at its baseline and is least similar to good 1 in this parameterization; when $\alpha = 1$, goods 1 and 2 are identical in characteristics and become perfect substitutes. Increasing α moves good 2 continuously toward good 1 in characteristics space.

We assume that the government levies a specific excise tax τ on good 1. The consumer problem is

$$\max_x u(C(\alpha)x) \ s.t. \ (p_1+\tau)x_1 + p_2 x_2 = m, \tag{1}$$

where p_i is the producer price of good i, τ is the tax on good 1, and m is the income of the agent.[4] We assume that the consumer prices of the two goods are equal initially (i.e. $p_1 + \tau = p_2$) as this simplifies the analysis somewhat. However, this is not necessary to derive qualitatively similar results.

2.1.2 Substitutability and demand elasticities

Our main result is that the relationship between substitutability and elasticity is highly non-linear and convex (see Appendix C for derivation):

$$\varepsilon_1(\alpha) \approx \frac{1}{1-\alpha}\varepsilon_1(0), \tag{2}$$

where $\varepsilon_1(\alpha)$ is the price elasticity of good 1 as a function of α and $\varepsilon_1(0)$ is the price elasticity of good 1 at $\alpha = 0$.

Figure 1 provides the intuition in the two-good model. It considers three starting points for α: low (left panel), medium (middle panel), and high (right panel). In each case, increasing substitutability by the same amount $\Delta\alpha$ straightens the indifference curve

[4]This highlights the connection between a Gorman–Lancaster model and the standard model of consumer choice: if there are as many goods as characteristics and C is the identity matrix, we have the standard model.

(black to red), making the consumer more willing to switch across goods when relative prices change.

The figure also links this geometry to tax responses. For a given tax change, the movement from x_0 to x_1 is larger when α is higher. If α is increased by $\Delta\alpha$, the new response is $x_1(\Delta\alpha)$, and the gap between $x_1(\Delta\alpha)$ and x_1 is much larger when the goods are already close substitutes. Thus, equal increments in α generate much larger changes in elasticity at high initial α than at low or medium α. This is the convex pattern summarized by Equation (2) and illustrated in Figure 2.

2.2 Welfare implications of sin taxes

This section characterizes the welfare effects of increasing a targeted consumption tax within the characteristics-based framework from Subsection 2.1. We consider a second-best environment in which sin taxes also raise revenue, so welfare effects combine a Pigouvian correction channel and a Ramsey revenue channel. In the main text we focus on the two-good case where only good 1 generates an externality or internality[5], and evaluate marginal tax changes at current tax rates.

2.2.1 Setup

We retain the two-good setting and notation from Section 2.1, but relax the equal-price normalization to allow for asymmetric taxation of good 1. For consumer j, the problem is

$$\max_{x_j} u_j(Cx_j) \quad \text{s.t.} \quad (p_1 + \tau)x_{1j} + p_2 x_{2j} = m_j,$$

where p_1 and p_2 denote producer prices and m_j is income.

We assume that the consumption of good 1 generates a negative externality. We model this through a welfare-loss function $e_j(c_{11}x_{1j})$, with $e_j'(\cdot) > 0$, where c_{11} scales the harm-relevant content delivered by good 1 (for example, sugar content).[6] Good 2 does

[5]We provide an n-good model and extensions in Appendix D.

[6]We consider more general externalities in Appendix D, but the main intuition stays the same. One

not generate an externality.

We assume a continuum of consumers of mass one. Let aggregate consumption of good 1 be $x_1 = \int_j x_{1j}$, and let ε_1 denote the aggregate own-price elasticity of demand for good 1, defined as the consumption-weighted average of individual elasticities.

The government values consumer welfare and faces an external revenue requirement R. Social welfare is given by

$$W(\tau) = \int_j w_j(\tau) - \int_j e_j(c_{11}x_{1j}(\tau)) + \mu\left[\tau x_1(\tau) + T - R\right], \tag{3}$$

where $w_j(\tau)$ denotes normalized indirect utility[7] and μ is the marginal value of public funds.

The externality term is separable from private utility, so social harm from good 1 is treated as an additional social cost. We consider non-separable externalities in Appendix D.

2.2.2 Ramsey and Pigouvian Components

Differentiating welfare with respect to the tax rate and re-arranging yields the expression:

$$\frac{dW}{d\tau} = \underbrace{-\frac{e'c_{11}}{p_1+\tau}\varepsilon_1 x_1}_{\text{Pigouvian component}} + \underbrace{-x_1 + \mu\left[x_1 + \frac{\tau}{p_1+\tau}\varepsilon_1 x_1\right]}_{\text{Ramsey component}}. \tag{4}$$

where e' denotes the consumption-weighted marginal external damage from good 1. The derivation follows standard envelope arguments and is provided in Appendix D. The Pigouvian component captures corrective gains from reducing harmful consumption, while the Ramsey component captures the revenue-efficiency trade-off under distortionary taxation.

Pure Ramsey case ($e' = 0$). Without externalities, welfare effects are purely fiscal: the tax is attractive when demand is inelastic and the base is stable, and costly when

could also interpret e_j as an internality if consumers do not fully account for future private costs.

[7]Dividing by the marginal utility of income: $w_j(\tau) = u_j(\tau)/\lambda_j(\tau)$, where $\lambda_j(\tau)$ is the marginal utility of income of individual j.

demand is elastic and behavioral distortions are large. In our framework this corresponds to low substitutability (low α, low $|\varepsilon_1|$), where consumers cannot easily avoid the tax.

Pure Pigouvian case ($\mu = 1$, no revenue motive). Without a revenue motive, welfare gains depend on how strongly the tax reduces harmful consumption, so both the marginal externality $e'c_{11}$ and the elasticity ε_1 matter. External harm alone is not enough: the tax must induce behavior change.

To see this, note that when $\mu = 1$, the welfare change in Equation (4) becomes:

$$\frac{dW}{d\tau} = \frac{\tau - e'c_{11}}{p_1 + \tau}\varepsilon_1 x_1.$$

So, even at the Pigouvian benchmark $\tau^* = e'c_{11}$, welfare gains are limited when ε_1 is close to zero and are larger when a close substitute exists (high α leading to high elasticity), because then consumers avoid the tax by consuming more good 2 that does not create the externality.

Optimal tax. For completeness, setting $dW/d\tau = 0$ in Equation (4) and solving for the optimal tax rate yields the standard result:

$$\frac{\tau^*}{p_1 + \tau^*} = \underbrace{\frac{1}{\mu}\frac{e'c_{11}}{p_1 + \tau^*}}_{\text{Pigouvian}} + \underbrace{\frac{1-\mu}{\mu}\frac{1}{\varepsilon_1}}_{\text{Ramsey}}. \tag{5}$$

The optimal ad valorem rate is the sum of a Pigouvian term and a Ramsey inverse-elasticity term (see e.g. Sandmo (1975)). When $\mu = 1$, the Ramsey term vanishes and $\tau^* = e'c_{11}$. Thus, the optimal tax depends on the elasticity only through the Ramsey term, although more broadly (non-optimal tax) the welfare considerations have that as discussed above.

2.2.3 Marginal Value of Public Funds

We summarize the welfare effects of a tax increase with the marginal value of public funds (MVPF), defined as welfare impact per euro of net fiscal cost for a marginal tax change (see Hendren (2016); Hendren and Sprung-Keyser (2020)). Note that the MVPF

calculation does not assume that taxes are at optimal levels.

In our setting, the MVPF takes the form:

$$\text{MVPF} = \frac{1 + \dfrac{e'c_{11}}{p_1 + \tau}\varepsilon_1(\alpha)}{1 + \dfrac{\tau}{p_1 + \tau}\varepsilon_1(\alpha)}. \tag{6}$$

The elasticity $\varepsilon_1(\alpha)$ enters both numerator and denominator, so its magnitude determines how informative the MVPF is for tax design.

Small elasticity: Ramsey logic dominates. When demand is inelastic (low α in our setting such that very close substitutes are not available), behavioral responses are limited, the corrective channel is weak, and the MVPF remains close to one. Taxation then mostly reflects revenue-raising logic: the tax base is stable, but corrective gains are small.

Large elasticity: Pigouvian logic can dominate. When demand is elastic (high α such that very close non-taxed and non-externality creating substitutes are available), behavioral responses are large and the MVPF becomes sensitive to whether τ is below or above the marginal external damage $e'c_{11}$. If $\tau < e'c_{11}$, corrective benefits dominate and MVPF > 1; if $\tau > e'c_{11}$, tax-base erosion dominates and MVPF < 1. Thus, when elasticity is large, getting the tax rate right matters much more, because the welfare consequences of over- or under-taxing are amplified by the strong behavioral response.

Figure 3 summarizes these cases. Panel (a) varies elasticity at fixed externality: higher elasticity steepens the MVPF response to tax misalignment. Panel (b) varies the externality at fixed elasticity: higher externality shifts the curves rightward, with each crossing MVPF $= 1$ at $\tau = e'c_{11}$. Together, the panels show that externalities determine the threshold, while elasticities determine the strength of welfare gains or losses around that threshold.

In summary, the MVPF analysis shows that sin taxes can increase welfare when the good that creates externalities or internalities have close enough substitute available that

consumers react to the tax by shifting consumption to it. Moreover, the substitute good should not create externalities (or create less of them).

3 Institutional description of the sin tax scheme

We study multiple reforms in a sin tax scheme called the "sweets tax" scheme in Finland. On 1 January 2011, Finland introduced an excise tax on sweets, ice creams and raised the existing excise tax on soft drinks.[8] There were two subsequent excise tax increases in 2012 and in 2014, while the tax was abolished for sweets and ice cream on 1 January 2017 and continued to be applied to soft drinks.

The sweets tax is based on custom categories. As a result, some products are taxable while others, quite similar ones, are not. For example, cookies, sweet pastries or snack bars are not tax liable while chocolates and candy bars are. Within a category the tax was assigned based on weight or volume of the package. The tax was, thus, not based on the sugar content of the product. Moreover, since the tax is tied to the weight or the volume, the tax impact relative to before tax price of a product varies across the products within a taxed product category.[9]

Figure 4 illustrates and Table 1 describes the multiple tax rate changes in the sweets tax scheme. There was an excise tax rate on soft drinks of 0.045 €/l before the introduction of the sweets tax. Since 1 January 2011, the sweets tax rates amounted to 0.75 €/kg for the solid taxable products and 0.075 €/l for the liquid taxable products including soft drinks. The rates were subsequently increased so that, from 1 January 2012 onward, the valid rates were 0.95 €/kg for solid taxable products and 0.11 €/l for liquid taxable products.

Importantly, the tax rates were also changed on 1 January 2014 so that the tax rate for the liquids including soda with a sugar content higher than 0.5 g per 100 g or 100

[8]The tax applied also to chewing gums, and some other naturally or artificially sweetened products. These products under sweets tax are also subject to the reduced VAT rate (13% from July 2010 until December 2012 and 14% since January 2013). The VAT is calculated based on the consumer price that includes the excise tax.

[9]Additionally, products used as intermediate inputs in production and small scale production remained untaxed.

Table 1: Tax changes in 2010-2017

		Solid	Liquid	
Old excise tax	Before 2011	0 €/kg	0.045 €/l	
Sweets tax introduced	1 Jan 2011	0.75 €/kg	0.075 €/l	
Rate increase	1 Jan 2012	0.95 €/kg	0.11 €/l	
			Sugar free	Sugared
Rate differentiation	1 Jan 2014	unchanged	unchanged	0.22 €/l
Sweets tax scrapped; old excise tax remains	1 Jan 2017	0 €/kg	unchanged	unchanged

As of 1 January 2014, the liquids with a sugar content higher than 0.5 g per 100 g or 100 ml were defined as sugared. The definition of the sugar-free liquids was based on the regulation (EC) No 1924/2006 of the European Parliament and of the Council of 20 December 2006 on nutrition and health claims made on foods.

ml was raised to 0.22 €/l. The tax rate for liquids having lower sugar content than the threshold remained at 0.11 €/l. The 2014 change resulted in very similar products having a different excise tax rate. The purpose of this differentiated tax change was to channel consumption into sugar-free products.

The definition of the sugar-free liquids was based on an existing legal definition for sugar-free products.[10] The use of an existing regulation and a very clear limit for sugar content (below the limit there are only waters and artificially sweetened products) decreased the administrative burden for firms and the tax authorities monitoring the scheme.

The sweets tax is carried out whenever products are provided for the consumption purposes from a tax-free warehouse or imported to Finland. Manufacturers, wholesale sellers and importers are the parties liable to tax. However, it is possible to export products from a tax-free warehouse without incurring tax liabilities. The Finnish Customs Office is responsible for carrying out the sweets tax collection and overseeing the producers.

[10]Regulation (EC) No 1924/2006 of the European Parliament and of the Council of 20 December 2006 on nutrition and health claims made on foods.

4 Data

4.1 Scanner data

Our data source is a large Finnish retail-chain called the S-group, and the data consist of their product-specific sales figures of all relevant food products at the outlet level aggregated to a weekly level from the last week of 2009 to the end of 2018. These outlets are evenly located all across the country and there are 1,077 of them. The total data consists of ca. 500 million observations. The market share of the S-group in the Finnish grocery market is roughly 45%, making the sales data quite representative of the consumption behavior, although it is not a random sample of stores in Finland. The S-group has various types of stores: they include small corner shops, larger market stores as well as big mall-type stores located in bigger cities and also throughout the country including in smaller municipalities. Thus, S-group sales data represent purchases in various types of stores and demographies.

The data consists of taxed and non-taxed product groups. In the taxed group we have: sweets (chocolates, candy bars and assorted candies, chewing gums), ice creams, drinks (soft drinks, juices, juice concentrates, waters). In the non-taxed group we have: cookies, cakes, sweet and salty snack bars or pastries, sugars, honeys, syrups, and yogurts.

The data include the product name (e.g. Angry Birds Red Bird), its custom category (e.g. 2106), its detailed product category (e.g. fruit xylitol chewing gums with a package size larger than 60 g), its net consumer package size[11] in kilograms or liters (e.g. 0.07 g) and whether this refers to liters or kilograms. The data also include the total number of the items sold and the monetary value of the items sold as well as information on the number of the discounted items and the total monetary values of these discounts. All the product-specific information is summed up for each outlet for each week. Finally, we have data on the type of store and which region it is located in.

All the monetary data is in euros. There are no data on the wholesale prices or the identity of the producers or consumers. Based on the product names, the product

[11]The definition refers to the package size net of packaging and non-edible ingredients, making it equivalent to the legal definition of the excise tax base.

categories and the custom categories, we identify which products are taxable and which are not.

Table 1 describes the data. The first three columns give mean standard deviation and number of observations for the taxed products, and the following three columns the same for the non-taxed products. For these prodicts we describe prices (tax inclusice) and quantities in units sold. Panel A describes the whole sample, panel B describes the values in 2010, before the 2011 reform, panel C the values for soft drinks in 2013 before the 2014 reform for them and panel D the values for ice creams and sweets in 2016 before the tax was abolished for them in the 2017 reform.

The table shows that on average the price of the taxable products is below 2 euros per unit and they are sold sold on average around 12 times per week in a store. Especially the latter has a large standard deviation indicating heterogeneity in the store size. The non-taxed products have roughly similar prices and quantities sold than the taxed products.

4.2 Survey and additional product data

In addition to the sales data, we conducted a survey eliciting preferences over substitutability across products. We also explain below the use of additional product-level data in order to construct measures of substitutability between product groups.

The survey was administered both as a class assignment (voluntary or compulsory, depending on the course) on several economics courses in Finnish universities and through an open online link on the home pages of Finnish Centre of Excellence in Tax Systems Research and the Labour Institute for Economic Research. The main body of the survey consists of questions about products presented in pairs. This part of the survey operates by showing a sequence of pairs of products in random order, and asking each time the same three questions about these pairs of products. Thus, one respondent answers the same questions from multiple pairs of products giving us greater statistical power. An example of a product pair in the survey with the related questions is presented in Figure A.1 in Appendix A. The product pairs are randomly chosen from a large pool of pairs of products.

An important part of the survey is a photo for each product, because these give essential information about the products. In addition to photos, the survey describes the names and prices of products. In order to get photos for all products we in some instances got a permission to use the official marketing photo for each product. Because we did not get the official photos for all products, we also set up a photography studio with studio lights to take photos that are of similar quality than the marketing photos.

The main question in the survey is Question 2 we use to construct a substitution index. The question asks how substitutable product A is with product B on a scale between 0 to 10. Respondents could use a slider bar to provide an answer or type in the number to a box next to the slider bar. The higher the number the more substitutable the products are with each other, which was also each time explained in the survey next to the question. We then average the answers for each product pair to derive a substitution index across product pairs.

In addition, Question 3 is a money metric question on the substitutability across products. It states that product A costs certain amount of euros, and what would the price of product B need to be in order you to substitute that product with product A. The answers to this question are averaged for a product pair to receive a money metric substitution index. A small complication with Question 3 is that the absolute differences in the value of products are not equivalent to relative differences as the products have different base values.

In total, there were 3346 respondents to the survey, with more than 35,000 product pairs viewed in total.

We also gathered additional information on products for a subgroup of our sales data. These data were scraped from the Foodie.fi webpage of the S-group with their permission using the EAN codes of the products. These data include information on the ingredients, nutritional content, and store-floor locations. We have data on ingredients for 6,546 products, nutrient data for 6,052 products, and store-floor locations (in a single typical hyper market) for 2,091 products. We derive substitution indices based on these metrics such that products that have similar ingredients are more substitutable.

4.3 Results for substitution preferences and substitutability

For the survey, Figure 5 shows the distributions of answers to Question 2 asking on a scale between 0 to 10 how substitutable products are with higher number indicating more substitutable. In the figure, we group the answers from this questions for five product categories facing tax increases in the reforms versus their closest category not facing tax changes at the same reform, that is, substitution product category. The closest substitution categories are cookies for sweets, ice-cream and chocolate, and sugar-free soda for sugary soda. The figure shows the distribution of how substitutable respondents think the products in the tax-change category are with their closest substitution category. The figure shows that the distribution has concentration of mass close to low substitutability values for sweets, and the values start to increase for ice-cream and chocolate and reach a concentration of high substitutability numbers for sugary soda and juice. This indicates, for example, that according to our survey chocolate and cookies are clearly less substitutable with one another than sugary and sugar-free soda or juice, because only the latter ones have a large mass close to one in the distribution of answers to Question 2.

Figure 6 shows a substitution index for the different product-group pairs considered above. The index is constructed by simply taking the mean of the answers to Question 2 for which Figure 5 provides the distributions, as well as the 95% confidence intervals of the means. An interesting observation is that the substitution index seems to increase roughly linearly across the product-group pairs considered. The substitution index constructed from the survey is the main information we use in the analysis below and when fitting the theory model with the empirical setting.

Figure A.2 presents the alternative substitution indices along with the main substitution index as discussed above. The alternative substitution indices develop in the same manner as the survey-based measure we use in the main analysis.

5 Empirical results and empirical strategy

In this section we first provide our empirical results both on the effects of the reforms in the sin tax scheme on prices and quantities sold. We then describe how our empirical results align with a prediction from our theory framework when using the substitution index to link them.

5.1 Estimation methods

In terms of the estimation method, the main interests are to estimate price incidence as well as to what extent tax changes affect consumption of goods that face a tax change, and to what extent there is substitution to products that do not face a tax change. To achieve that, we examine in the event-study framework changes in the consumption of different product groups: those that face tax changes and those that are potentially close substitutes using more distant groups as the control group. We also quantify the effects by estimating differences-in-differences estimates, and provide elasticity of consumption with respect to prices estimates.

Relative to a simple quasi-experimental setting, where one would have a treated and control group, we have a more challenging problem of determining which are the potential substitution groups and which are the product groups that can be used as a control. Choosing the product groups affects also whether they are comparable over time and thus whether the control group establishes a credible counterfactual. Thus, we need to be careful in how we choose the groups. In choosing them, we relied on our results from Section 4.3 on the substitutability of different product groups and implemented the following procedure.

We first compared product groups facing tax reforms with various product groups that did not face tax changes. Out of these, we selected all products that seemed like close substitutes based on the survey described above to form a potential substitution group.

We then selected a control group for both the tax group and the substitution group

from products that are fairly similar to them, for example all products that can be used as snacks, to credibly represent a counterfactual, but that were not very close substitutes to the taxed or substitute group in the survey. In the latter case we selected only product groups for which we managed to establish parallel pre-trends with the taxed or substitute groups. Fortunately, our data included suitable products that both had parallel pre-trends and that were not very close substitutes, such as, snacks, yogurts and pastries.

This procedure was sufficient to provide an estimate for all products facing a tax change in one of the reforms in our setting.

We present the event-study graphs both such that different groups described above are shown separately or that their difference in each time period is shown as a normal event-study graph, to provide a transparent description of the results. We estimate the following equation:

$$y_{it} = \sum_{\tau=-T,\tau\neq-1}^{T} \beta_\tau [t = \tau] g_i + \gamma_t + \mu_i + \epsilon_{it} \tag{7}$$

where i refers to a product sold in particular outlet and τ refers to time periods (either week or month, depending on the setup) and g_i to the group identifier, either one facing tax change, potential substitute group or control group. The event-study specification is such that the g_i is interacted with event time τ, and when we estimate just the time-series for the groups separately we do not include this interaction. μ_i refers to product-outlet i fixed effects and γ_t refers to a set of time fixed effects. ϵ_{it} is the residual error term. The outcome variables y_{it} consist of volume of each product sold or their log price (euros). The standard errors are clustered at the product-outlet level.

We impute the mechanical full tax pass-through to graphs examining price-incidence. A factor here is that the data consists of retail prices that include the excise tax and VAT. We impute the base price for each product as its mean price in November and the two first weeks of Decembers of 2010, 2013 and 2016 before each reform. Based on this product-specific base price, we impute the full pass-through as follows and take the

logarithm over it:

$$\ln\left(p_{full,it}\right) = \ln\left(\left(1 + \text{VAT }_t\right)\left(\frac{\text{base price}_{it}}{1 + \text{VAT }_t} + \text{new ex. tax }_{it}\text{-old ex. tax}_{it}\right)\right) \quad (8)$$

The logarithms of the product-specific pass-through rates are then averaged over for the relevant product group (i.e. all the products, solids only etc.). The full pass-through of the tax is calculated by adding this mechanical price effect due to taxes to the non-taxed group prices.

When analyzing quantities sold, one additional methodological concern is that the sales of the products in our sample exhibit strong seasonal fluctuation as exemplified in Figure 7. The figure compares the development of the amount of taxable products (excluding soft drinks) around the 2011 and 2012 reforms (mainly sweets, chocolate and ice cream) with a set of non-taxable products (mainly cookies, sweet and savoury pastry, pudding and yogurt). The extent of seasonality is high in that sales vary 20% around the annual average sales. We consider seasonality as something naturally arising from consumption varying with season, for example, influenced by seasonal temperature. One example is that ice cream is consumed much less in the winter than in the summer in Finland. A concern is that not taking seasonality into account would cause a bias in estimations. In order to adjust for the seasonality, we use the residuals from an estimation that regresses the respective quantity variable on the product category-calendar month effect.[12]

We quantify the impact of the sweets tax reforms with a DiD framework. We estimate the following equation:

$$y_{it} = \gamma_t + \mu_i + \beta_{DD} a_t \times g_i + \epsilon_{it}, \quad (9)$$

where a_t is an indicator taking value one for after the reform period, g_i is the same group indicator as above, $a_t * g_i$ is the interaction of the two indicator terms and thus β_{DD} is the

[12]The seasonality highlights the value of having observations for each month over many years over having only few isolated points in time, for example, from a survey, because then producing de-seasonalizing the data would be impossible.

main coefficient of interest. y_{it} is the outcome variable, either unit price or the volumes or items sold. We include in the estimation also product-outlet level fixed effects μ_i. ϵ_{it} is the error term, which is clustered at the product-outlet level.

We also provide an estimate for the consumption elasticities. We do this in two-stage least squares, where we first estimate the impact of a reform on log prices and use this variation to explain the change in log quantities.

$$\begin{aligned} \log p_{it} &= \gamma_t + \mu_i + \beta_{DD} a_t \times g_i + \epsilon_{it}, \\ \log q_{it} &= \gamma_t + \mu_i + \log\hat{p}_{it} + \varepsilon_{it} \end{aligned} \tag{10}$$

Where we first estimate the impact of the reform on log prices $\log p_{it}$ and in the second equation explain the change in log quantities $\log q_{it}$ with the estimated change in log prices $\log\hat{p}_{it}$. Both equations include the product by store fixed effects μ_i as well as time fixed effects γ_t .

5.2 Graphical evidence on the effects of the sin tax scheme on prices and quantities

Here we provide graphical evidence on how reforms impacted prices and quantities sold of products by estimating equation (7) and following the procedure explained in the previous subsection.

We first document the evidence on the impact of the reforms in the sin tax scheme on prices. Figure 9 presents the development of prices at weekly level around the 2011 and 2012 reforms for all taxed and non-taxed products pooled. The imputed full pass-through is marked as the dashed black line. The figure shows that the prices of taxed, potential substitution group (Cookies) and non-taxed products follow on average the same trend in the year before the reform. Then at the time of the 2011 reform the taxed group prices exhibit a clear increase. The prices do not jump immediately to the new level in 2011, but rather the increase is gradual over the period of three months ending up to a level of full pass-through on average. In the 2012 reform prices jump more rapidly to an

increased level. The prices end up at a higher level after 2012 than what the mechanical full-pass through would indicate. The price of Cookies do not show such marked increases but ends up at a slightly higher level than the non-taxed group prices, which could be just natural price development and not affected by taxes to the taxed group that are not applied to Cookies.

After the two reforms in 2011 and 2012 the taxed group is at around 10% higher level than the control group, indicating the effect of the two reforms combined on prices on average.

To provide some evidence for different subcategories, Figure **??** presents the development of prices separately for ice creams, soft drinks, sweets and chocolates (as a special category of sweets) in the 2011 reform. Relative to the average price effect of the two reforms, ice cream prices are different because they increase by less than the amount of the mechanical full pass-through. This arises possibly due to the larger product weight-based excise taxes for this group relative to pre-reform prices. Instead, for soft drinks the pass-through was less than full, and for this group the tax increase due to the reform was smaller. Thus, there is some evidence of an inverse relationship between the relative size of taxes and the degree of pass-through.

Figure 10 presents the price incidence for abolishing the sweets tax for sweets and ice cream from the beginning of 2017. The figure indicates an immediate and large price response to the reform. However, on average the prices declined by less than what the full pass-through would have been. This suggests some modest asymmetric pass-through by comparison to the tax introduction, which is somewhat consistent with the findings of Benzarti et al. (2020) for VAT. The sub-group analysis in Figure **??** show that for sweets the pass-through seems to be full for this reform, but for the other categories slightly less.

For the 2014 reform differentiating the tax for sugary and non-sugary soft drinks, Figure 11 shows that the consumer prices responded immediately by increasing to a level of mechanical full pass-through after the reform. The figure also shows the prices of sugar-free liquids and the prices of a control group. The sugar-free drinks exhibit initially some price increase after the reform, but this subsides and their prices soon return to the

same leven with the control group. Importantly, the effect of the 2014 reform on prices of sugary drinks is about 10%, which is very similar to the average price effect in absolute terms of the other reforms.

In summary, the reforms had a clear bite on prices, such that each reform we analyze had an average effect of changing consumer prices by about 10% in absolute value. Thus, it is conceivable that if consumption elasticity is sizable, we should detect a change in the consumption of the targeted products proxied by the sales quantity response. Also, when comparing below the effects of the different reforms on quantities sold, any potential differences across reforms is not likely to be attributable to differential price effects of the reforms.

Turning to analyze the impact of sin taxes on quantities sold, Figure 13 presents the over time development of the amount of products sold for the 2011/2012 reforms. The series is de-seasonalized. The figure presents again three groups (g_i), those facing the tax change, which are sweets and ice cream here, cookies as the closest substitute group and a control group consisting of similar but more distant products, mainly sweet and savory pastries, puddings and yoghurt, in panel (a). In panel (b) the figure present event-study graph that compares the monthly level difference of taxable group relative to the control group. The figure shows that the three groups develop on the same trend before the 2011 reform. However, the groups develop similarly also after the two reforms in 2011 and 2012. Thus, the figure shows quite convincingly that the two reforms had no large impact on the sales of the taxable group, and no substitution effect to the closest substitution group.

Figure 13 presents the development of soft drinks, cookies as a potential substitution group and a control group comprising puddings in the 2011/2012 reform. The emerging picture is very similar to above, there does not seem to be a discernible effect on the sales of soft drinks in the 2011/2012 reform.

Figure 14 presents the development of sales for the three groups over the 2017 reform that abolished the tax for sweets and ice cream, but left the tax for soft drinks in place. The sales of the three groups presented in panel (a) develop again quite smoothly over

the 2017 reform without any clear and systematic differences between the groups after the reform. As for the introduction of the tax in 2011, this figure does not depict any clear and detectable changes in the sales of the taxed or potential substitution group for abolishing the tax in 2017.

For the 2014 reform, Figure 15 shows the development of sales of sugary and sugar-free drinks including all drinks that faced the tax reform in the taxable group, as well as a control group (comprising honey and salty licorice).[13] Figure 16 presents the same comparison focusing on sodas and Figure 17 for focusing on juices as subgroups of all soft drinks. The figures show that the three groups are on a parallel trend prior to the 2014 reform and that there is now a clear separation between the sales of the taxable and potential substitution groups precisely at the time of the reform. Relative to the control group the consumption of sugary soft drinks declined and the consumption of sugar-free soft drinks increased. Thus, the graphs point to markedly different impact from the 2014 reform than the impact of 2011 and 2017 reforms discussed above, including the result in Figure 13 for the same product group when the sugary and non-sugary drinks experienced the same tax impact in the 2011/2012 reforms.

Furthermore, to examine the role of substitution separated from an overall impact on soft drink sales, Figure 18 depicts the total soft drink sales (adding sugary and sugar-free soft drinks together) relative to the control group in the 2014 reform. Here we do not observe any decline in total soft drink sales, which suggests that the behavioral response was mainly substituting from sugary to sugar-free drinks.

5.3 The effects of different reforms on quantities, prices and the implied consumption tax elasticities

In this subsection, we provide differences-in-differences estimates of the impacts of the sweets tax on quantities sold, prices and the estimated consumption elasticities with respect to prices. In addition to standard direct effects of sin taxes on taxed goods, we

[13]The 2014 reform doubled the tax rate for sugared liquids from 0.11 €/l to 0.22 €/l, while it left the tax on sugar-free drinks unchanged at 0.11 €/l.

want to identify substitution effects on goods that are not directly taxed but are potential substitutes based on our survey, as the methods section above explained.

Table 2 collects our differences-in-differences (DiD) estimates (equation (9)) for log sales and prices for the 2011 and 2012 reforms combined, 2017 reform and 2014 reform. For these reforms the table shows the DiD estimates for the taxed group and the closest substitution group relative to the relevant control groups. For the 2011 reform the pre-period is 2010 and after-period 2011 to 2013. For the 2014 and 2017 reforms we include one year before and one year after the reforms in the estimates. The outcome variable is log price or log volume of a product sold within a month in a store. We include product by store fixed effects in the estimation.

The DiD estimate quantifies as the price impacts on the taxed groups 15% increase in prices in the 2011/2012 reforms, -9% decline in prices in the 2017 reform and 9% increase in the 2014 reform. The price effects on the potential substitution groups are otherwise very close to zero, but cookies prices are estimated to increase by 2.6% in the 2011/2012 reforms. For quantities the DiD estimates are very close to zero for the 2011/2012 and 2017 reforms for the taxed and potential substitute groups being between -0.2% and 0.6%. Instead, for the 2014 reform the DiD estimate for the tax change group (sugary soft drinks) is -4.1% and for the substitute group (non sugary soft-drinks) 6.5%.

Table 3 collects our preferred quantity, price and elasticity estimates for different product categories facing tax changes, that also have varying products as the closest substitutes not facing tax changes. Here we use a combined effect from the 2011/2012 and 2017 reforms for sweets, ice-cream and chocolate. The table shows a pattern where the elasticity estimates are close to zero for sweets (0.022), ice-cream (-0.031) and chocolate (-0.08), while the estimate is markedly higher for sugary soda (-0.53) and sugary juice (-0.58). The main difference across these product groups is how close substitute the closest potential substitution group is. We investigate how these observation support the results from our theory framework in the next subsection.

5.4 Fitting the empirical estimates with our theory framework

We have presented above in Section 2 a theoretical framework according to which we organize the discussion of empirical estimates. Here we present how the empirical elasticity estimates align with the theory framework matched with the substitution index from our survey. The substitution index and how it is constructed based on survey answers is described in subsection 4.3. We use the substitution index to order the elasticity estimates according to their distance to closest substitution group that does not face a tax change. We also use the substitution index to draw a prediction from our model indicating what kind of pattern of consumption elasticities the theory predicts with the given substitution index. We then examine whether the pattern of quasi-experimental estimates organized in the same way as the prediction from the theory model align with each other.

In Figure 19, we present both the prediction from the theory framework and our quasi-experimental estimates in a two-dimensional graph, together with the substitution indices based on the EU tariff tree. On the horizontal axis of this graph is our Substitution index from the survey between the tax change group and the product group that is the closest substitute to it. On the vertical axis of the graph is the consumption elasticity.

We fit a version of Equation 2 to the data on elasticity estimates and substitution indices using non-linear regression with the functional form

$$\varepsilon_i = \frac{1}{1 - \beta_0 SI_i}\beta_1, \tag{11}$$

where ε_i is the estimated elasticity of a reformed product category (e.g. chocolate, or sugary soda), and SI_i the survey based substitution index between the reformed product category and it's closest substitute category. Here we assume that $\alpha_i = \beta_0 SI_i$. Moreover, β_1 estimates the baseline elasticity $\varepsilon(\alpha = 0)$.

Figure 19 presents the resulting relationship between the substitution index both from our theory model shown by the shaded gray area and our quasi-experimental elasticity estimates for different product groups marked by the blue dots. The solid gray line shows the main estimated prediction from the theory model and the shaded area corresponds

to varying the parameter values of the model. As the gray line in the figure shows, the theory predicts that the demand elasticity of the taxed products increases in a highly non-linear fashion with the substitution index. The blue dots in the figure showing the quasi-experimental estimates align well with the model prediction. We estimate a demand elasticity of close to zero for the other pairs than soft drinks and juices. The effect is dramatically different for sub-groups of soft drinks and juices for which we estimate high demand elasticities and which also lie on the high end of our substitution index.

6 A meta-analysis on previous literature and substitutability

To further describe the relationship between substitutability to non-taxed goods and consumption elasticities, we conducted a meta-analysis of the elasticity estimates in the literature and related these to an empirical measure of substitution possibilities based on EU tariff-categories.

We included in our meta-analysis papers that provide quasi-experimental consumption elasticity estimates with respect to prices utilizing changes in excise taxes and sales tax or VAT on various products. We collected 29 different estimates from 22 studies. The studies are reported in Table 4.

6.1 Assigning substitutability measures to previous literature

To understand the relationship between price elasticity estimates from previous literature and the substitutability possibilities underlying those elasticities, we build a simple index measuring substitution possibilities in each case. The index is based on the EU tariff goods-tree, which assigns categories to product groups in an descending order of aggregation and thus produces a tree branch type categorization. With the goods-tree, we measure substitution possibilities as a function of the shortest distance to a product group for which tax treatment does not change, and hence forms a potential substitution group.

We use the goods-tree with 3-levels of categories with the finest level being 6 digits categories.[14] The main rule we use to construct the index is as follows: for a given good in a given setting, we obtain a tariff category code at the 6-digit level (for example, sugary soda in 042202). Next, using the goods-tree, we find a closest good available which taxes do not change (e.g. artificially sweetened soda, which is in the same category 042202). This translates to four possible cases: 1) the closest substitute may be in the same category, 2) the closest substitute may be in a neighboring category, so that the four leading figures of the code are the same (e.g. for sweetened soda in 042202, non-sweetened soda is in category 042201), 3) the closest substitute shares only the first two digits (e.g. for sweetened soda in category 042202, this is a candy in category 041704 in the absence of other sodas), or 4) the closest substitute does not share any digits, because non exist within the same branch of the product tree. Finally, for each of these cases, we assign an index number between 0 and 1 as follows: we compute the shortest distance from a taxed good to a non-taxed good in the tariff tree (e.g. when soda sweetened with sugar or artificially in 042202 is taxed, one has to go through two intersections connecting different branches to get to non-sweetened soda in 042201) this distance is then normalized to be between 0 and 1 by dividing by the longest possible route (with three levels, this means six intersections, so 2/6 for sweetened soda). Next, we turn this normalized distance (d) into a measure of substitution possibilities $s = 1 - d^2$. The idea here is that the number of goods increases approximately to the power of two with distance[15] An example of this is illustrated in Figure 7.

However, not all goods and no services are covered by the EU tariff goods-tree, as they are not accounted for by the tariff authorities. Within our reviewed elasticity estimates, such goods are electricity, restaurant meals, hair-cuts at hairdressers, and miles/kilometers travelled by a car. We assign all these to the category with the least

[14]For example, for sugary soda, the code is 042202, where 04 stands for "Prepared foodstuffs; beverages, spirits, and vinegar; tobacco and manufactured tobacco substitutes", 22 stands for "Beverages, spirits and vinegar", and 02 stands for "Waters, including mineral waters and aerated waters, containing added sugar or other sweetening matter or flavored, and other non-alcoholic beverages, not including fruit, nut or vegetable juices of heading 2009".

[15]Using two is only an approximation and could in practice be a higher number as there tend to be more than two sub categories within each higher level category.

substitution possibilities (index of 1/6) based on the following reasoning: Electricity usage at the household-level could be substituted mostly in the case of heating by other heating sources (e.g. natural gas), but this requires a renovation at the energy system of the house and takes time. Other household electricity consumption may be harder to find a substitute for. Restaurant meals could perhaps be substituted by home-prepared meals, but restaurants tend to be used for different types of occasions (e.g. parties or work lunches), and restaurant meals may ask for expert skills and equipment that one might not have available for home-prepared meals. Similar reasoning applies to hair-cuts, which might be substituted for by home production, but which may demand expert skill and equipment not available at home. Miles/kilometers of travel by car might be substituted with other forms of travel, but in many cases such other forms of travel may be nonexistent.

6.2 Elasticity estimates as a function of substitution possibilities in the previous literature

Using the procedure described above we categorize the studies in the meta-analysis into four substitution groups. The substitution index allocated to each group is the average of the substitution indices allocated to each study in the group. The blue dots in Figure 21 show the average elasticities organized by the average substitution indices in the horizontal axis. The capped bars show the average confidence intervals using the standard errors from the studies in the meta-analysis. The figure shows that the relationship between the average consumption elasticities and the substitution index has a convex shape. Thus, the result aligns well with the prediction from our theory framework, which is presented with a stylized grey curve behind the average elasticities from the literature.

Table 4 gives the details of the studies used. Different columns in the table give the product-group that is taxed in the study, the closest substitute product-group, which one of the four substitution groups the estimate falls to, the substitution index we derived for the substitution group on average, and the consumption elasticity estimated in the study.

As the table describes, in the highest substitution group are studies on taxing sugar-sweetened beverages with artificially flavored products as the closest substitution group and where the reform occurs in a narrow geographical area. Also in the highest substitution group are studies on taxing cigarettes or e-cigarettes with other tobacco products as the closest substitution group. In the second highest substitution group are studies where the taxed group is various alcohol products with possibility to import alcohol from neighboring countries, or studies on sweetened beverages where the tax is applied to a wide geographical area. In the third group are studies on taxing gasoline, electricity or mile of travel by car, the we deem have as the closest substitution group other forms of transport. In the lowest substitution group are studies on value added taxes or a broad tax on all sweets and beverages in Norway that has as the closest substitution products other groceries.

We performed robustness checks and varied some of the choices we made in the analysis discussed above. These are reported in the Appendix in Figures A.3 through A.5. The result is that the convex pattern of elasticities still emerge even if we exclude the studies where we had to make some assumptions about the substitution index or make alternative assumptions about how to calculate standard errors.

7 Conclusions

We study under what conditions sin taxes and other targeted consumption taxes increase welfare. The idea is that much of responses to such consumption taxes might be driven by substitution towards goods that are taxed at lower rates. To more formally describe this mechanism, we present a theory framework that allows us to derive analytical results for how demand elasticity depends on how similar, or substitutable, taxed and non-taxed goods are. The model provides as a prediction that the relationship with the subsitutitability and demand elasticity is a highly convex shape, such that the demand elasticity is high only for goods that have very close non-taxed substitutes available.

The size of the elasticity is connected with the welfare implications. The welfare

increases more when taxing goods that have very close substitutes available, because then the consumer response to taxes is greater. Another aspect affecting welfare is the externalities (or internalities) the consumption of goods create in relation to each other. Welfare increases the most for goods that create larger externalities and that have as close substitutes goods that create no or much less externalities.

We study these aspects empirically by utilizing a sin tax scheme in Finland with multiple reforms. Although there is a substantial impact to the consumer prices, we did not find a sizeable impact on the sales of goods from tax reforms that treated very close substitute groups similarly. In contrast, we find that the 2014 tax increase for sugary soft drinks that left the tax of a very close substitute group of non-sugary soft drinks unchanged did have a large impact on the sales of sugary drinks explained by substitution.

For price-incidence we see that all reforms led to a roughly 10% change in prices on average. Thus, differential price impact across different reforms cannot explain the very differential impacts on consumption. One interesting finding is that the price-incidence is modestly asymmetric; abolishing the tax was under shifted to prices while tax increases were fully or even over shifted to prices. For example, Benzarti et al. (2020) find more significant asymmetry for VAT, and the modest asymmetry here could be related to the different nature of the excise tax levied on producers relative to the VAT on consumption goods.

In order to link the theory model with the empirical results we provide a substitution index describing how close substitutes the products facing tax changes are with products not facing tax changes. Our main index is based on a survey that we designed for this purpose. The survey indicates that sugary and non-sugary soft drinks are the most substitutable with each other from different taxed groups, while for other taxed product groups the substitution index is slightly lower.

We find that the prediction from our theory model aligns well with our empirical estimates, when ordering them based on the substitution index. Thus, we have empirical support for the theory model that predicts the size of the demand elasticity for sin taxes and other targeted consumption taxes based on the availability of close substitutes.

We also provide a meta-analysis to the literature providing consumption elasticities estimated in quasi-experimental settings. We develop a substitution index for these studies based on EU tariff-goods tree and augmented with geographical proximity and other considerations of proximity for services not found in the tree. We organize the earlier studies into four categories based on the substitution index and take the average of the estimated elasticities within group. We find again a convex relationship such that the large demand elasticities are found only in the highest substitution category while other average elasticity estimates are much closer to zero. Thus, the earlier literature on average aligns with our substitution based theory prediction.

The substitution indices developed in this study are useful for ex ante policy considerations of what would happen following a particular targeted consumption tax reform, when combined with our theory framework. The framework also allows to perform ex ante welfare calculations based on the assumed size of externalities, elasticities based on the substitution possibilities and the existing level of taxes.

References

(2010): "Laki virvoitusjuomaverosta 17.12.2010/1127 (Law on soft drink taxes)," .

(2013): "Sokeriverotyöryhmän loppuraportti," .

Agrawal, D. R. and W. H. Hoyt (2019): "Pass-Through in a Multiproduct World," *Working paper*.

Allcott, H., B. B. Lockwood, and D. Taubinsky (2019a): "Regressive Sin Taxes with and Application to the Optimal Soda Tax," *The Quarterly journal of economics*, 134, 1557–1626.

——— (2019b): "Should We Tax Sugar-Sweetened Beverages? An Overview of Theory and Evidence," *The Journal of economic perspectives*, 33, 202–227.

Allcott, H. and C. Rafkin (2022): "Optimal Regulation of E-cigarettes: Theory and Evidence," *American economic journal. Economic policy*, 14, 1–50.

Andreyeva, T., M. W. Long, and K. D. Brownell (2010): "The Impact of Food Prices on Consumption: A Systematic Review of Research on the Price Elasticity of Demand for Food," *American journal of public health (1971)*, 100, 216–222.

Asplund, M., R. Friberg, and F. Wilander (2007): "Demand and distance: Evidence on cross-border shopping," *Journal of public economics*, 91, 141–157.

Benzarti, Y., D. Carloni, J. Harju, and T. Kosonen (2020): "What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes," *The Journal of political economy*, 128, 4438–4474.

Berardi, N., P. Sevestre, M. Tépaut, and A. Vigneron (2016): "The impact of a 'soda tax' on prices: evidence from French micro data," *Applied economics*, 48, 3976–3994.

Bergman, U. M. and N. L. Hansen (2019): "Are excise taxes on beverages fully passed through to prices? The danish evidence," *Finanzarchiv*, 75, 323–356.

Cawley, J. (2015): "An economy of scales: A selective review of obesity's economic causes, consequences, and solutions," *Journal of health economics*, 43, 244–268.

Cawley, J., D. Frisvold, and D. Jones (2020): "The impact of sugar-sweetened beverage taxes on purchases: Evidence from four city-level taxes in the United States," *Health economics*, 29, 1289–1306.

Cawley, J. and D. E. Frisvold (2017): "The Pass-Through of Taxes on Sugar-Sweetened Beverages to Retail Prices: The Case of Berkeley, California," *Journal of policy analysis and management*, 36, 303–326.

Cawley, J., A. M. Thow, K. Wen, and D. Frisvold (2019): "The Economics of Taxes on Sugar-Sweetened Beverages: A Review of the Effects on Prices, Sales, Cross-Border Shopping, and Consumption," *Annual review of nutrition*, 39, 317–338.

Chan, T. Y. (2006): "Estimating a continuous hedonic-choice model with an application to demand for soft drinks," *The Rand journal of economics*, 37, 466–482.

Chetty, R., A. Looney, and K. Kroft (2009): "Salience and Taxation: Theory and Evidence," *The American economic review*, 99, 1145–1177.

Chiou, L. and E. Muehlegger (2014): "CONSUMER RESPONSE TO CIGARETTE EXCISE TAX CHANGES," *National tax journal*, 67, 621–650.

COGLIANESE, J., L. W. DAVIS, L. KILIAN, and J. H. STOCK (2017): "ANTICIPATION, TAX AVOIDANCE, AND THE PRICE ELASTICITY OF GASOLINE DEMAND," *Journal of applied econometrics (Chichester, England)*, 32, 1–15.

Colantuoni, F. and C. Rojas (2015): "THE IMPACT OF SODA SALES TAXES ON CONSUMPTION: EVIDENCE FROM SCANNER DATA: IMPACT OF SODA SALES TAXES," *Contemporary economic policy*, 33, 714–734.

Collaborators, T. G. . O. (2017): "Health Effects of Overweight and Obesity in 195 Countries over 25 Years," *The New England journal of medicine*, 377, 13–27.

Cotti, C., C. Courtemanche, J. C. Maclean, E. Nesson, M. F. Pesko, and N. W. Tefft (2022): "The effects of e-cigarette taxes on e-cigarette prices and tobacco product sales: Evidence from retail panel data," *Journal of health economics*, 86, 102676–102676.

Cotti, C., E. Nesson, and N. Tefft (2016): "The Effects of Tobacco Control Policies on Tobacco Products, Tar, and Nicotine Purchases among Adults: Evidence from Household Panel Data," *American economic journal. Economic policy*, 8, 103–123.

Dubois, P., R. Griffith, and A. Nevo (2014): "Do Prices and Attributes Explain International Differences in Food Purchases?" *The American economic review*, 104, 832–867.

Dubois, P., R. Griffith, and M. O'Connell (2020): "How Well Targeted Are Soda Taxes?" *The American economic review*, 110, 3661–3704.

Evans, W. N., J. S. Ringel, and D. Stech (1999): "Tobacco taxes and public policy to discourage smoking," *Tax Policy and the Economy*, 13, 1–55.

Fletcher, J. M., D. E. Frisvold, and N. Tefft (2010): "The effects of soft drink taxes on child and adolescent consumption and weight outcomes," *Journal of public economics*, 94, 967–974.

Gillingham, K., A. Jenn, and I. M. Azevedo (2015): "Heterogeneity in the response to gasoline prices: Evidence from Pennsylvania and implications for the rebound effect," *Energy economics*, 52, S41–S52.

Gillitzer, C., H. J. Kleven, and J. Slemrod (2017): "A Characteristics Approach to Optimal Taxation: Line Drawing and Tax-Driven Product Innovation," *The Scandinavian journal of economics*, 119, 240–267.

Goetzke, F. and C. Vance (2021): "An increasing gasoline price elasticity in the United States?" *Energy economics*, 95, 104982.

Gonçalves, J. and J. P. dos Santos (2020): "Brown sugar, how come you taste so good? The impact of a soda tax on prices and consumption," *Social science medicine (1982)*, 264, 113332–113332.

Gorman, W. M. (1980): "A Possible Procedure for Analysing Quality Differentials in the Egg Market," *The Review of economic studies*, 47, 843–856.

Grogger, J. (2017): "Soda Taxes and the Prices of Sodas and Other Drinks: Evidence from Mexico," *American journal of agricultural economics*, 99, 481–498.

Gruber, J., A. Sen, and M. Stabile (2003): "Estimating price elasticities when there is smuggling: the sensitivity of smoking to price in Canada," *Journal of health economics*, 22, 821–842.

Harberger, A. C. (1962): "The incidence of the corporation income tax," *Journal of Political Economy*, 70, 215–240.

Harding, M., E. Leibtag, and M. F. Lovenheim (2012): "The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data," *American economic journal. Economic policy*, 4, 169–198.

Harju, J. and T. Kosonen (2010): "Ravintoloiden arvonlisäveroalennuksen hintavaikutukset," *VATT-muistio 10*.

Harju, J., T. Kosonen, and O. N. Skans (2018): "Firm types, price-setting strategies, and consumption-tax incidence," *Journal of public economics*, 165, 48–72.

Harju, J., T. Kosonen, and J. Slemrod (2020): "Missing miles: Evasion responses to car taxes," *Journal of public economics*, 181, 104108.

Hendren, N. (2016): "The Policy Elasticity," *Tax Policy and the Economy*, 30, 51–89.

Hendren, N. and B. Sprung-Keyser (2020): "A Unified Analysis of Government Policies," *Quarterly Journal of Economics*, 135, 1209–1318.

Hindriks, J. and V. Serse (2022): "The incidence of VAT reforms in electricity markets: Evidence from Belgium," *International journal of industrial organization*, 80, 102809.

Hines, J. R. (1999): "Three Sides of Harberger Triangles," *The Journal of economic perspectives*, 13, 167–188.

Johansson, P., T. Pekkarinen, and J. Verho (2014): "Cross-border health and productivity effects of alcohol policies," *Journal of health economics*, 36, 125–136.

Kleven, H. and J. Slemrod (2009): "A Characteristics Approach to Optimal Taxation and Tax-Driven Product Innovation," *Unpublished manuscript*.

Kosonen, T. (2015): "More and cheaper haircuts after VAT cut? On the efficiency and incidence of service sector consumption taxes," *Journal of public economics*, 131, 87–100.

Kosonen, T. and J. Harju (2013): "Restaurant VAT cut: Cheaper meal and more service?" *IDEAS Working Paper Series from RePEc*.

Kosonen, T. and O. Ropponen (2012): "Makeisvero – tehokasta kulutusverotusta vai kulutuskäyttäytymisen ohjausta?" *VATT-muistio 21*.

Kotakorpi, K., T. Härkänen, P. Pietinen, H. Reinivuo, I. Suoniemi, and J. Pirttilä (2011): "Terveysperusteisen elintarvikeverotuksen vaikutukset kansalaisten terveydentilaan ja terveyseroihin," *Terveyden ja hyvinvoinnin laitos (THL): Raportti 7/2011*.

Lancaster, K. J. (1966): "A New Approach to Consumer Theory," *The Journal of political economy*, 74, 132–157.

Langer, A., V. Maheshri, and C. Winston (2017): "From gallons to miles: A disaggregate analysis of automobile travel and externality taxes," *Journal of public economics*, 152, 34–46.

Levin, L., M. S. Lewis, and F. A. Wolak (2017): "High Frequency Evidence on the Demand for Gasoline," *American economic journal. Economic policy*, 9, 314–347.

Lovenheim, M. F. (2008): "How Far to the Border?: The Extent and Impact of Cross-Border Casual Cigarette Smuggling," *National tax journal*, 61, 7–33.

Merker, T. (2023): "Consumer Self-Sorting and the Effect of Consumption Taxes," *Unpublished manuscript*.

Miravete, E. J., K. Seim, and J. Thurk (2020): "One Markup to Rule Them All: Taxation by Liquor Pricing Regulation," *American economic journal. Microeconomics*, 12, 1–41.

O'Donoghue, T. and M. Rabin (2006): "Optimal sin taxes," *Journal of public economics*, 90, 1825–1849.

OECD (2020): "The OECD Tax Database," .

Pigou, A. C. (1920): *The Economics of Welfare*, London: Macmillan.

Powell, L. M. and J. Leider (2020): "The impact of Seattle's Sweetened Beverage Tax on beverage prices and volume sold," *Economics and human biology*, 37, 100856–100856.

——— (2021): "Impact of a sugar-sweetened beverage tax two-year post-tax implementation in Seattle, Washington, United States," *Journal of public health policy*, 42, 574–588.

Powell, L. M., J. Leider, and P. T. Léger (2020): "The impact of a sweetened beverage tax on beverage volume sold in cook county, Illinois, and its border area," *Annals of internal medicine*, 172, 390–397.

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

——— (2015): "A CONTRIBUTION TO THE THEORY OF TAXATION," *The Economic journal (London)*, 125, 254–268.

Roberto, C. A., H. G. Lawman, M. T. LeVasseur, N. Mitra, A. Peterhans, B. Herring, and S. N. Bleich (2019): "Association of a Beverage Tax on Sugar Sweetened and Artificially Sweetened Beverages With Changes in Beverage Prices and Sales ay Chain Retailers in a Large Urban Setting," *JAMA : the journal of the American Medical Association*, 321, 1799.

Rojas, C. and E. Wang (2021): "DO TAXES ON SODA AND SUGARY DRINKS WORK? SCANNER DATA EVIDENCE FROM BERKELEY AND WASHINGTON STATE," *Economic inquiry*, 59, 95–118.

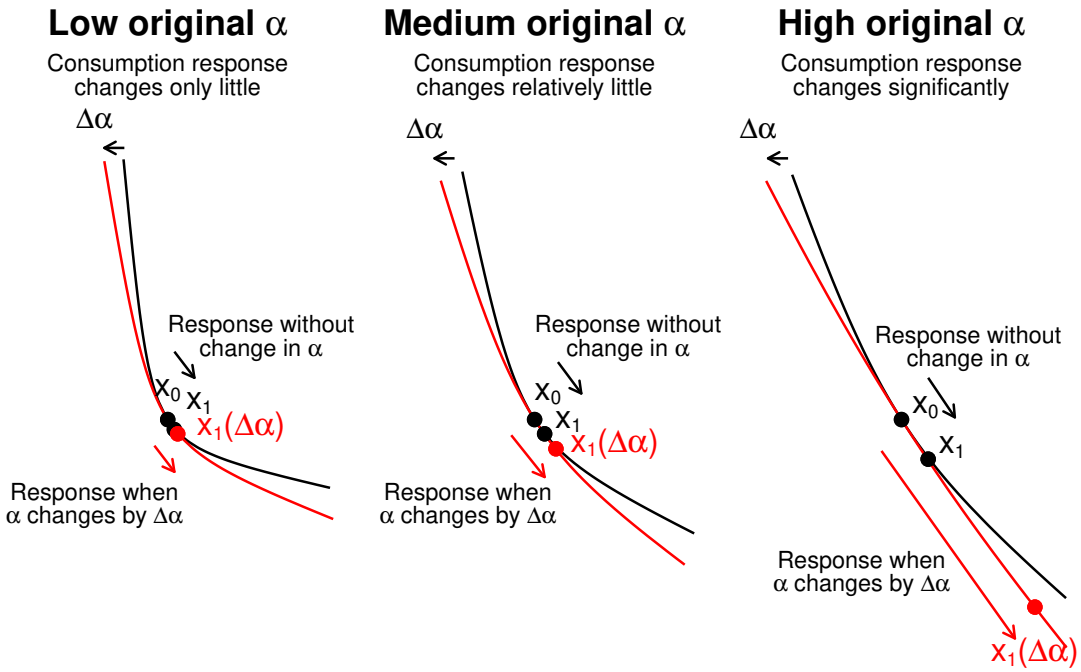
Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *The Journal of political economy*, 82, 34–55.

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

Seiler, S., A. Tuchman, and S. Yao (2021): "The Impact of Soda Taxes: Pass-Through, Tax Avoidance, and Nutritional Effects," *Journal of marketing research*, 58, 22–49.

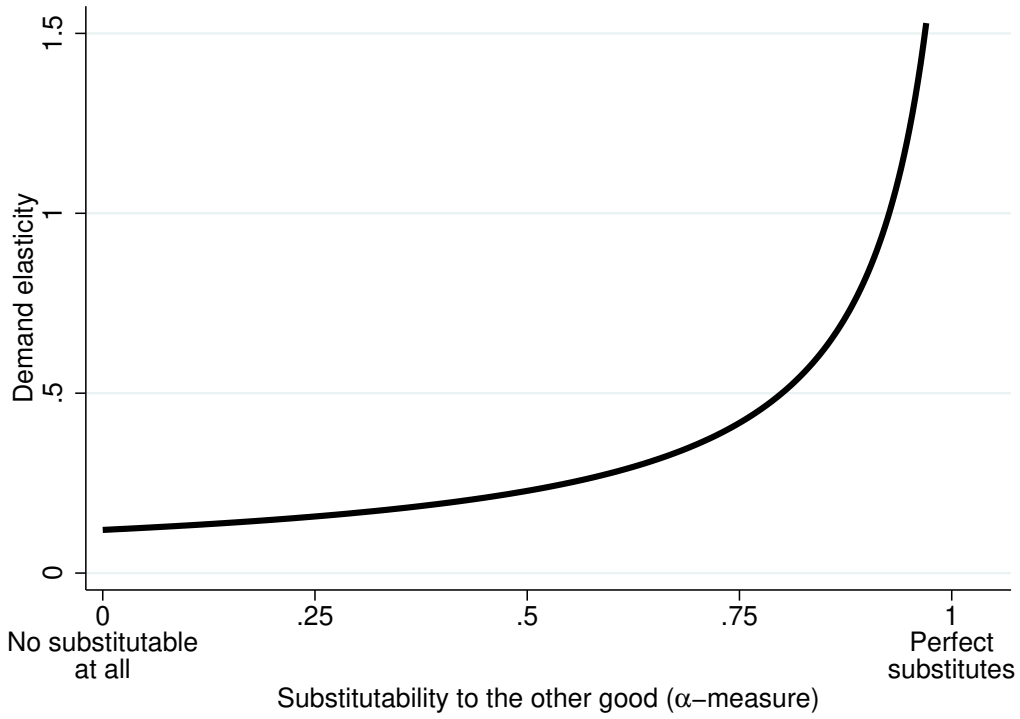
Figures

Figure 1: Impact of increased substitutability (α) on indifference curves in a stylized Gorman–Lancaster model



Note: The figure illustrates how increasing substitutability (labeled α) between goods in a stylized two-good Gorman–Lancaster (Gorman 1980, Lancaster 1966) model impacts the indifference curves between the two goods. This is measured by how similar the characteristics vectors of the goods are as outlined in section 2, so that when $\alpha = 0$, the goods are not substitutable at all, and when $\alpha = 1$, the goods are perfect substitutes. The impact of a small increase in substitutability ($\Delta\alpha$) on the curvature of the indifference curve is dependent on the original level of substitutability between the goods. The left panel depicts a situation with low original substitutability between the two goods (low α), whereas the middle panel depicts a situation with medium initial substitutability, and the right panel a situation with high initial substitutability. Increasing substitutability by $\Delta\alpha$ in all panels moves the indifference curves from the black indifference curves to the red ones. If we consider a tax change and the resulting change in the consumption bundle either from x_0 to x_1 (without $\Delta\alpha$) or from x_0 to $x_1(\Delta\alpha)$ (with the increase in α by $\Delta\alpha$), we can see, that when the goods are close substitutes to begin with (right panel), the impact is of the slight increase in substitutability ($\Delta\alpha$) induces a much larger change in the response to the tax change than when the goods are either not close (left panel) or intermediately close substitutes (middle panel). This implies a nonlinear, convex relationship between demand elasticities and substitutability between the two goods, as outlined in figure 2.

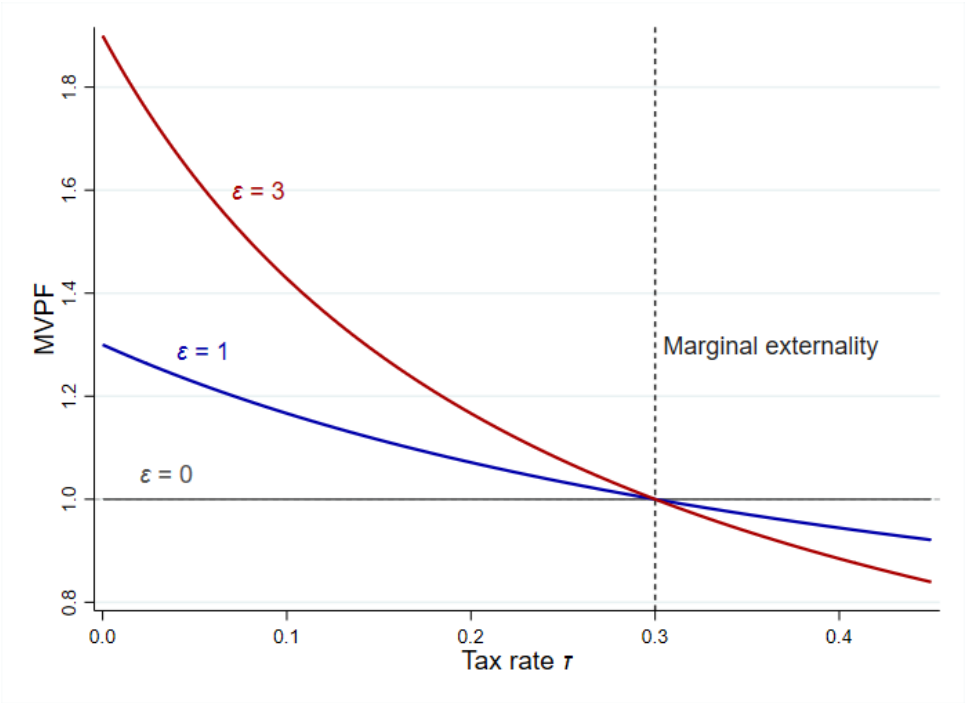
Figure 2: Demand elasticity as a function of substitutability to the other good in a stylized two-good Gorman – Lancaster model



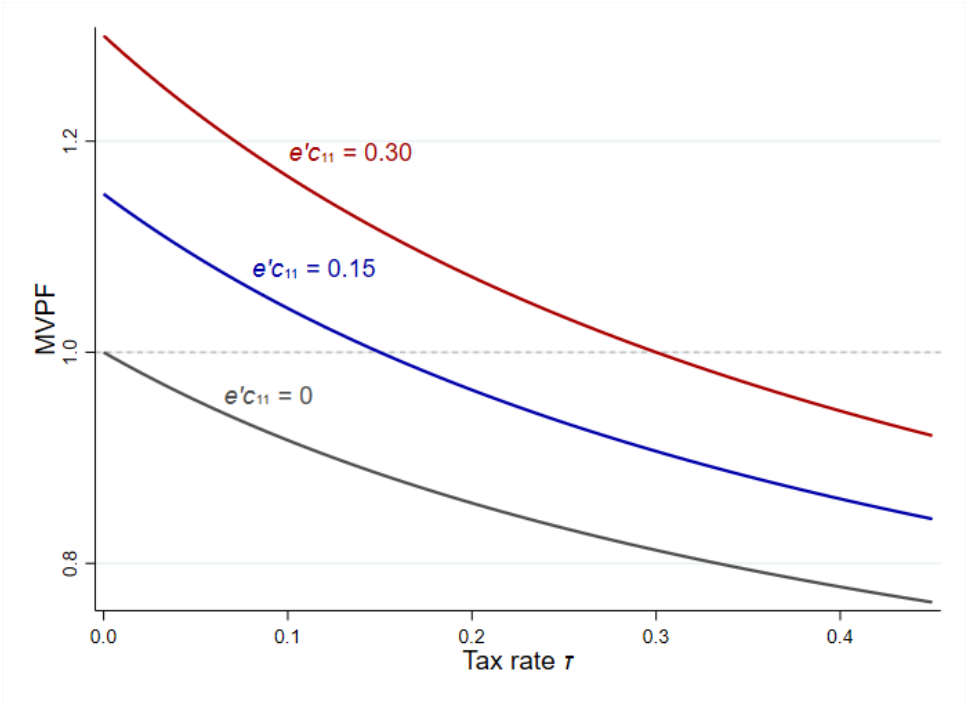
Note: The figure shows the (compensated) own-price demand elasticity of a good in a two-good, two-characteristics Gorman-Lancaster model as a function of how substitutable the goods are with each other. This is measured by how similar the characteristics vectors of the goods are as outlined in section 2, so that when $\alpha = 0$, the goods are not substitutable at all, and when $\alpha = 1$, the goods are perfect substitutes. Based on this stylized model, we would expect to find a convex relationship between substitutability and demand elasticities. This could potentially explain some of the large variation in elasticity estimates in different studies examining seemingly similar situations and goods.

Figure 3: MVPF as a function of the tax rate

(a) Different elasticities



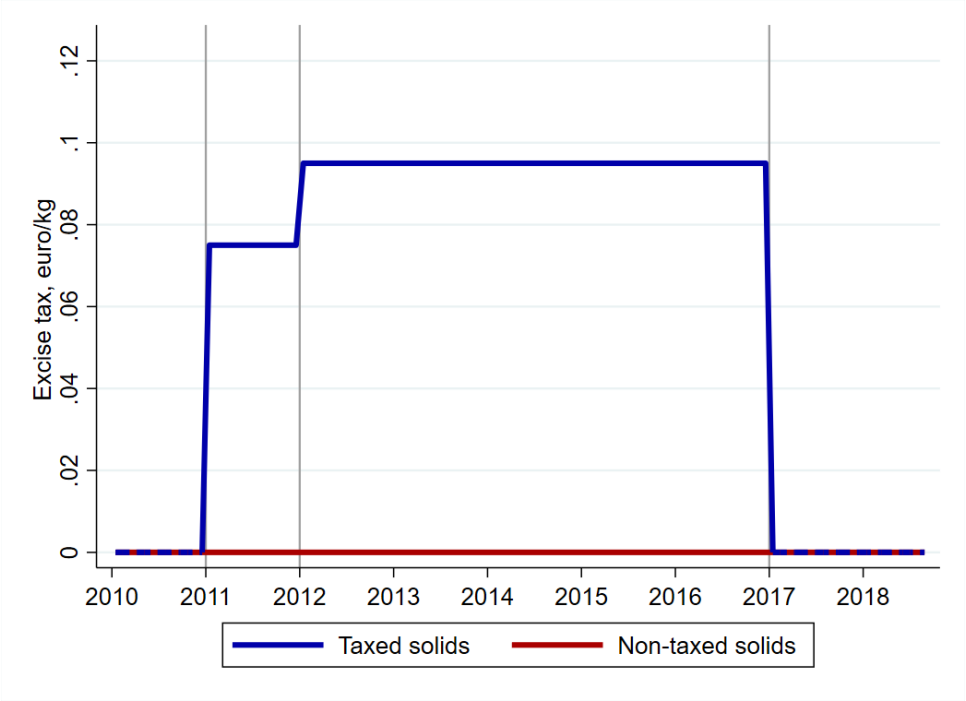
(b) Different marginal externalities



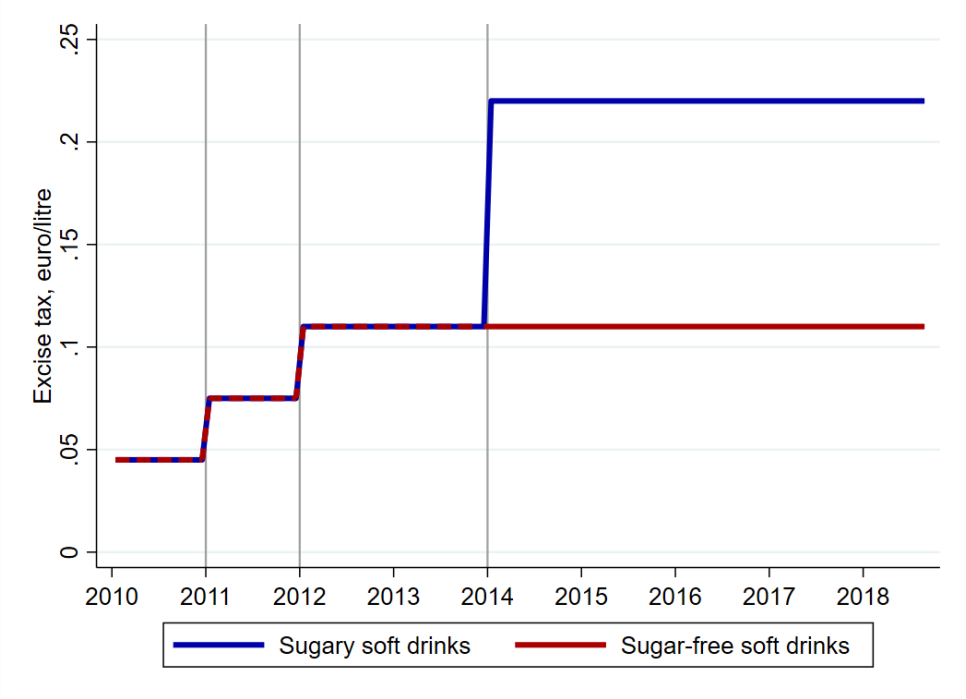
Note: The figure illustrates how the MVPF (from Equation (6)) of a tax change depends on the tax rate for different values of consumption elasticities and marginal externalities. Panel (a) shows three cases: elasticities of 0, 1 and 3. The marginal externality is held constant at 0.3. If the elasticity is 0, the MVPF is 1 at all times and the Ramsey logic dominates: one dollar of additional revenue from this tax reduces tax-payer welfare by one dollar, regardless of there being an externality. In contrast, if the elasticity is high, the Pigou logic dominates: the MVPF can be larger or smaller than one depending on whether the tax rate is larger or smaller than the marginal externality. This illustrates that the demand elasticity provides important information for welfare assessments of sin taxes. Panel (b) shows three cases: marginal externalities of 0, 0.15 and 0.3. The elasticity is held constant at 1. The curves have the same shape, but are shifted to the right with increasing marginal externality, with each crossing MVPF=1 at the level of the marginal externality.

Figure 4: Excise tax reforms for treated groups

(a) Taxed solids and control groups



(b) Sugary and sugar-free soft drinks and control groups

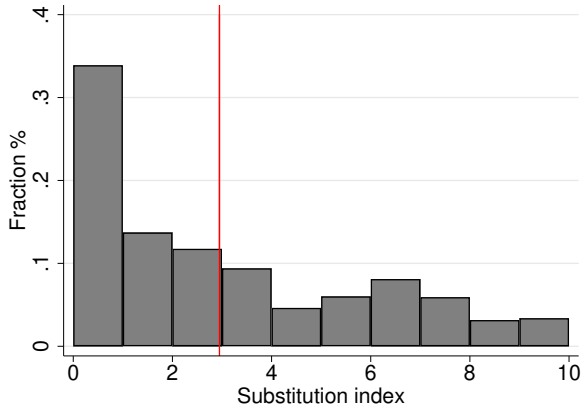


Note: Panel (a) shows how the excise tax changed for treated solids such as sweets, chocolate and ice-cream. Before 2011, the excise tax rate for these was 0 €/kg, but was raised to 0.075 €/kg in 2011 and to 0.095 €/kg in 2012. The tax for solids was abolished in 2017, so that the excise tax became 0 €/kg.

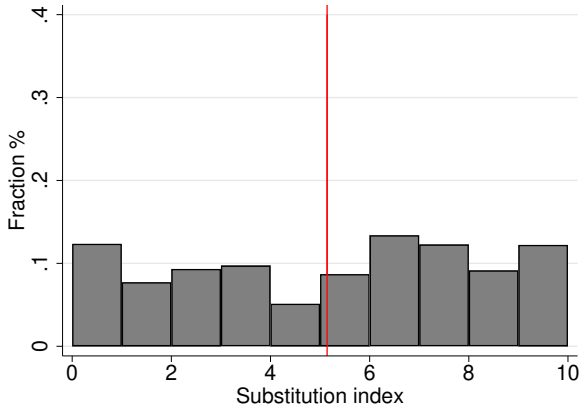
Panel (b) shows how the excise tax changed for treated soft drinks such as sodas and juices. Before 2011, the excise tax rate for these was 0.045 €/l, but was raised to 0.075 €/l in 2011 and to 0.11 €/l in 2012. In 2014 the tax for sugary soft drinks (more than 0.5 g of sugar per 100 g or 100 ml) was raised to 0.22 €/l, but the tax rate for sugar-free soft drinks stayed at 0.11 €/l.

Figure 5: Distributions of substitution preferences

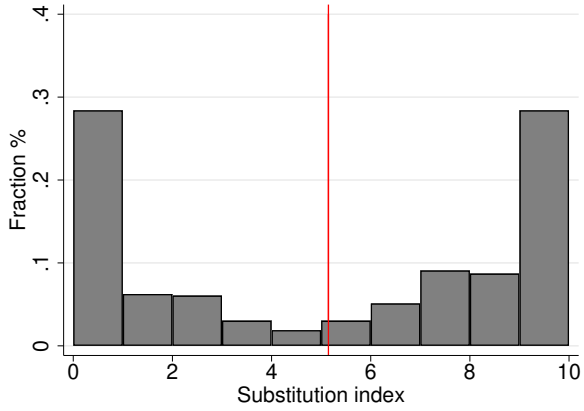
(a) Sweets vs. cookies



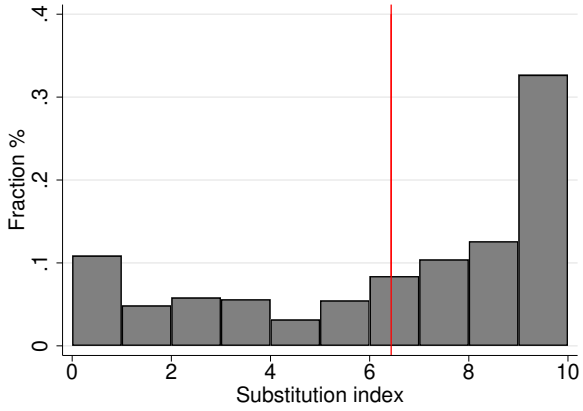
(b) Chocolate vs. cookies



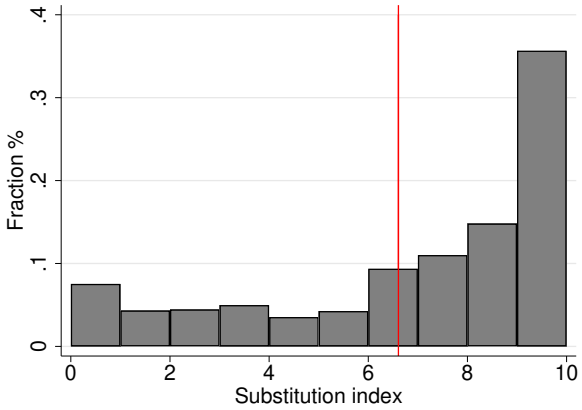
(c) Ice-cream vs. soda



(d) Sugary soda vs. sugar-free soda

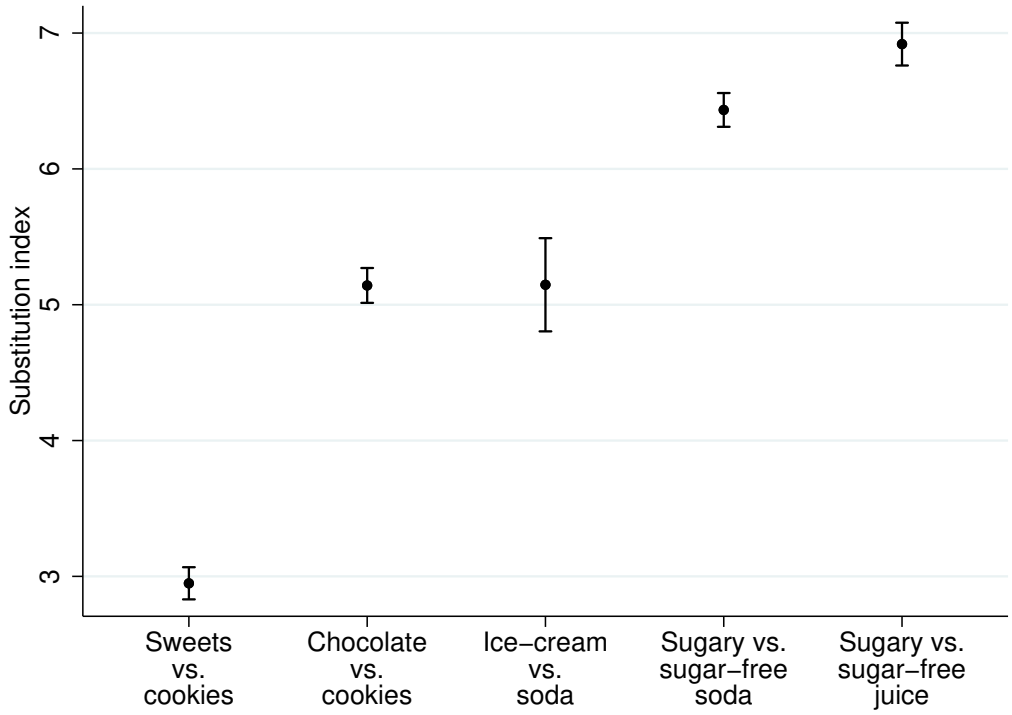


(e) Sugary juice vs. sugar-free juice



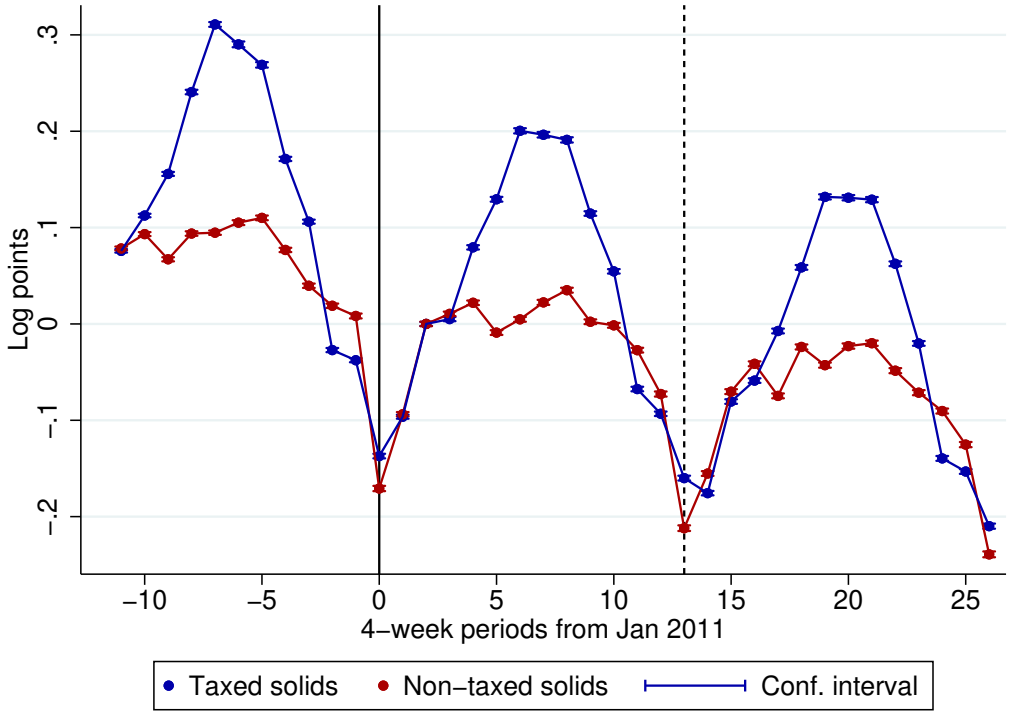
Note: The panels shows distributions of the answers to the survey question "On a scale from 0 to 10, how substitutable are these products for you?" for taxed product categories compared to their closest non-tax category. (Panel (a) Sweets vs. cookies, Panel (b) Chocolate vs. cookies, Panel (c) Ice-cream vs. soda, Panel (d) Sugary soda vs. sugar-free soda, Panel (e) Sugary juice vs. sugar-free juice.) Red vertical lines show the corresponding means. The closest non-tax category is chosen by the same index. The more the distribution is located to the right, the more substitutable the respondents have viewed the categories to be.

Figure 6: Means of substitution preferences



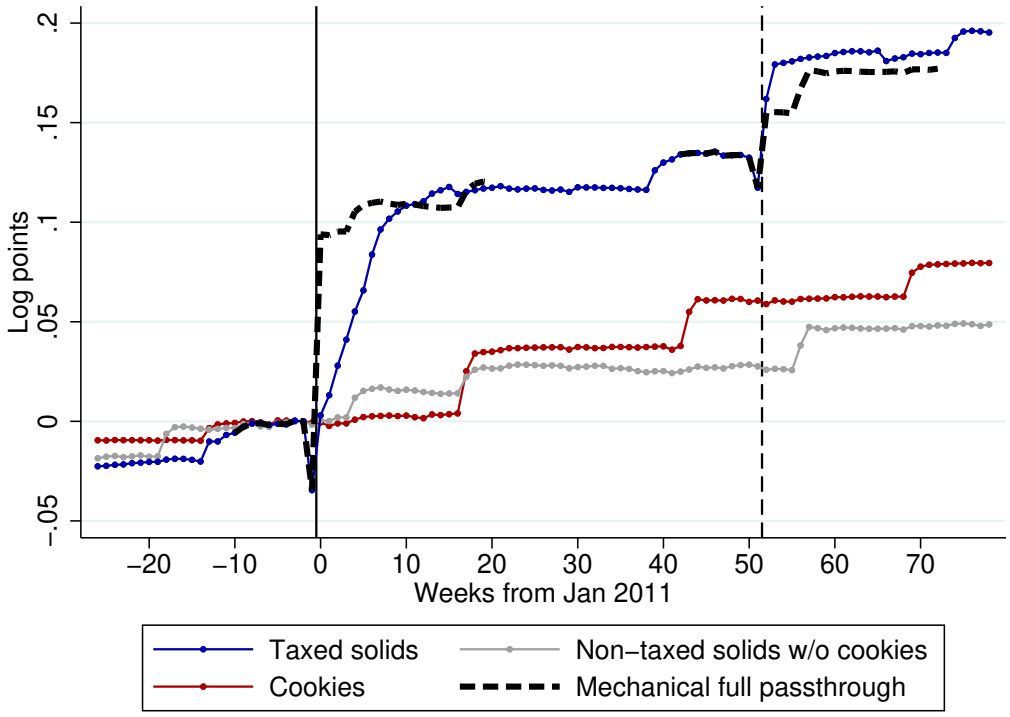
Note: The figure shows mean values of the answers to the survey question "On a scale from 0 to 10, how substitutable are these products for you?" for taxed product categories compared to their closest low-tax category. The closest low-tax category is chosen by the same index. The higher the value of the index, the more substitutable the respondents have viewed the categories to be. 95% confidence intervals are also included.

Figure 7: Development of quantities, solid products in the 2011 and 2012 reforms, not de-seasonalized



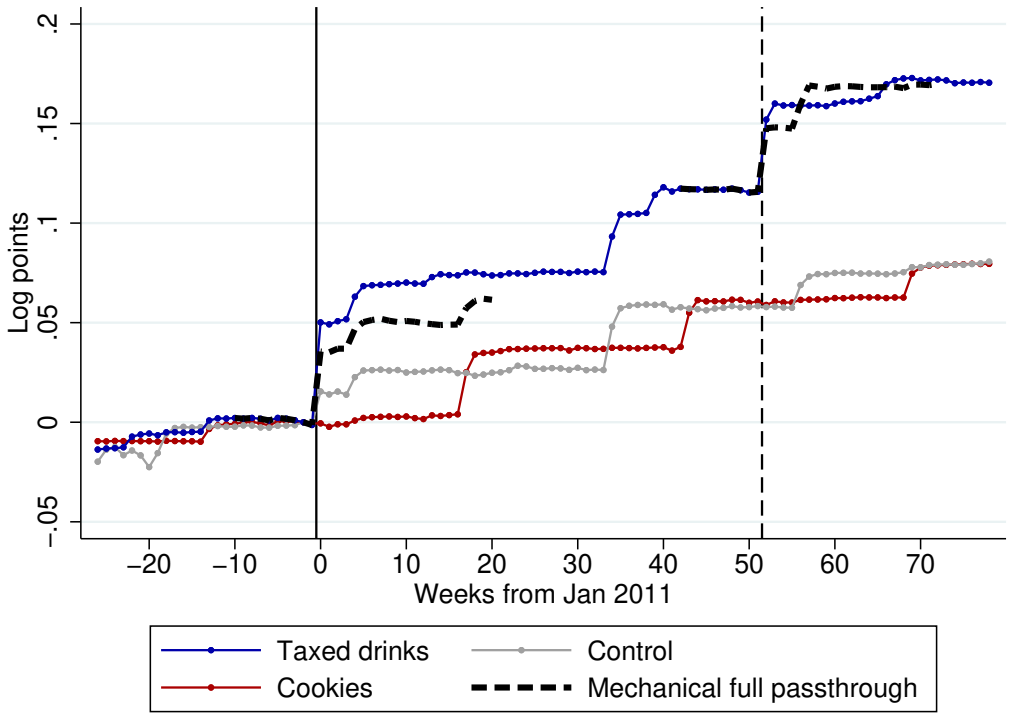
Note: The figure shows regression coefficients for month dummies from regressions of log number of products sold with product-outlet fixed effects (estimates of equation (7)). The graph displays strong within year seasonality, such that the difference between the peak in summer months to the lowest point in winter is close to 0.4 log points. Taxable products include all solid taxable products in the 2011 and 2012 reforms: mainly sweets, chocolate and ice cream. Non-taxable products include solid products that were not included in the 2011 and 2012 reforms: mainly cookies, sweet and savoury pastry, pudding and yogurt. The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level. Vertical lines indicate the 2011 sweets tax introduction and the 2012 tax increase. The series are normalized to zero at the last week prior to the 2011 reform.

Figure 8: Development of log unit prices of solid products in the 2011 and 2012 reforms



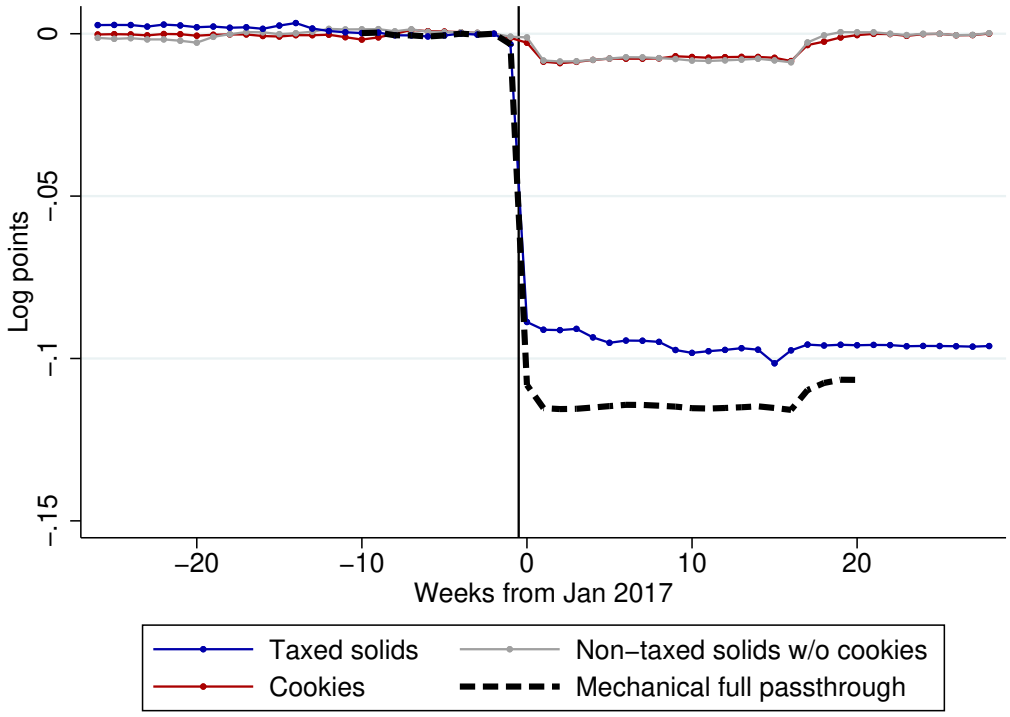
Note: The figure shows regression coefficients for week dummies from regressions of log unit prices with product-outlet fixed effects (estimates of equation (7)) for three groups: taxed solids, non-taxed solids without cookies, and cookies. Taxed solids include all taxed solid products in the 2011 and 2012 reforms: mainly sweets, chocolate, and ice cream. Non-taxed solids include all solid products in our data that were not included in the 2011 and 2012 reforms: mainly sweet and savory pastry, sweeteners, yogurt, and pudding. The figure shows that prices of taxed solids increase after the reform, and that there is a general increase in the prices of non-taxed solids as well. The prices of taxed solids end up at about 0.1 log points higher level than the prices of non-taxed solids after the 2012 reform. Mechanical full pass-through in dashed black line shows how prices of taxed solids would have developed with full pass-through of tax to prices, had they otherwise followed the prices of non-taxed products (wihout cookies) post-reform. Vertical lines indicate the 2011 and 2012 reforms. The series are normalized to zero at two weeks prior to the 2011 reform.

Figure 9: Development of log unit prices of drinks in the 2011 and 2012 reforms



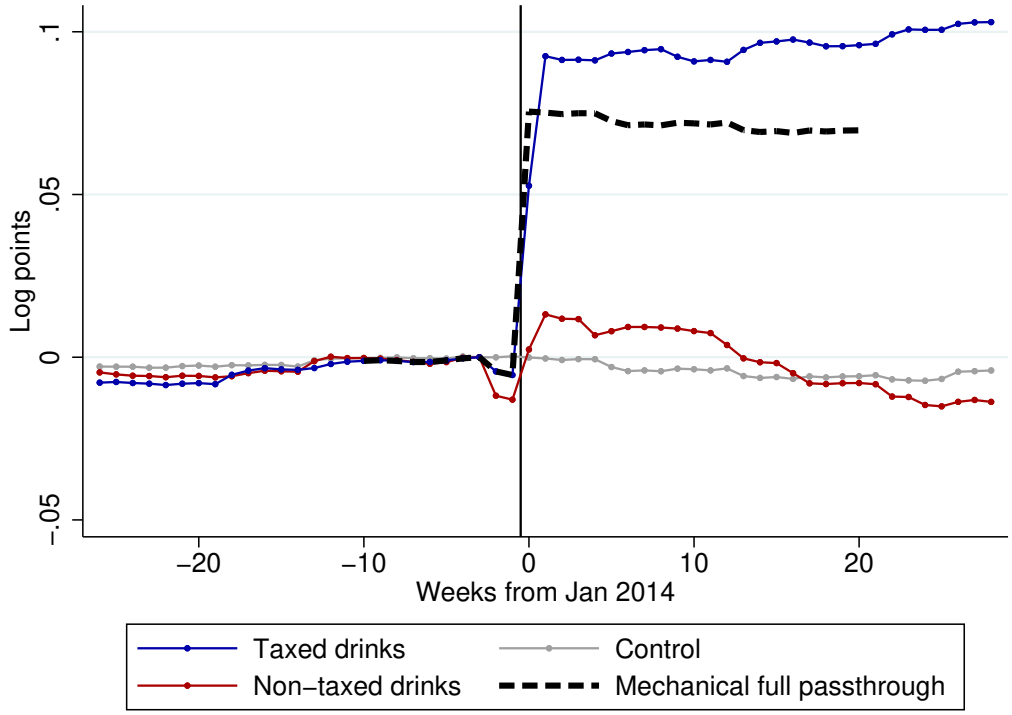
Note: The figure shows regression coefficients for week dummies from regressions of log unit prices with product-outlet fixed effects (estimates of equation (7)) for three groups: taxed drinks, a control group of puddings, and cookies. Taxed drinks include all taxed drinks in the 2011 and 2012 reforms: mainly soda and juice. The figure shows that prices of taxed drinks increase after the reform, and that there is a general increase in the prices of the other groups as well. The prices of taxed drinks end up at about 0.1 log points higher level than the prices of the control group after the 2012 reform. Mechanical full pass-through in dashed black line shows how prices of taxed drinks would have developed with full pass-through of tax to prices, had they otherwise followed the prices of the control group post-reform. Vertical lines indicate the 2011 and 2012 reforms. The series are normalized to zero at two weeks prior to the 2011 reform.

Figure 10: Development of log unit prices in the 2017 reform



Note: The figure shows regression coefficients for week dummies from regressions of log unit prices with product-outlet fixed effects (estimates of equation (7)) for three groups: taxed solids, non-taxed solids without cookies, and cookies. Taxed solids include all taxed solid products in the 2017 reform: mainly sweets, chocolate, and ice cream. Non-taxed solids include all solid products in our data that were not included in the 2017 reform: mainly sweet and savory pastry, sweeteners, yogurt, and pudding. The figure shows that prices of taxed solids decrease after the reform, and that there is otherwise no trend in the prices of non-taxed solids. The prices of taxed solids end up at about 0.1 log points lower level than the prices of non-taxed solids after the 2017 reform. Mechanical full pass-through in dashed black line shows how prices of taxed solids would have developed with full pass-through of tax to prices, had they otherwise followed the prices of non-taxed products (wihout cookies) post-reform. Vertical line indicates the 2017 reform. The series are normalized to zero at two weeks prior to the 2017 reform.

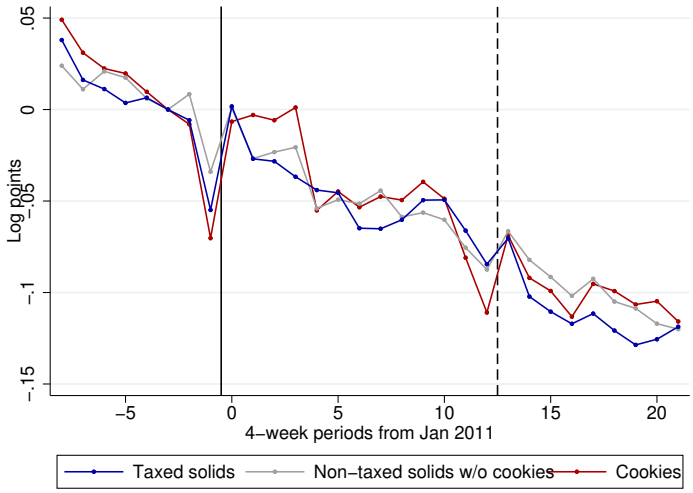
Figure 11: Development of log drink prices by sugar content and comparison prices in the 2014 reform



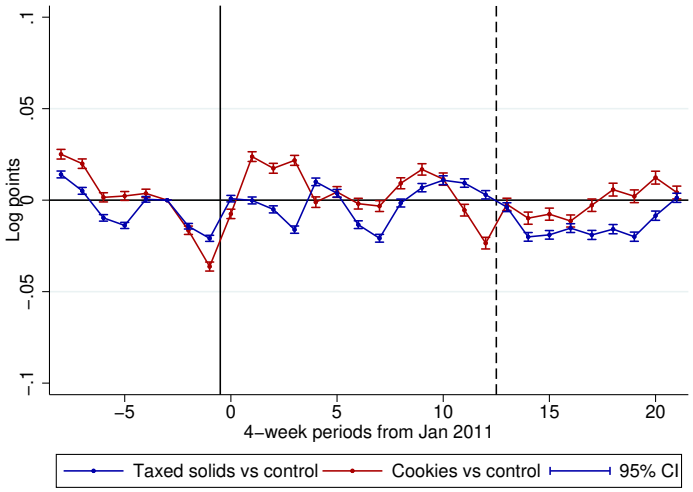
Note: The figure shows regression coefficients for week dummies from regressions of log volume-based prices with product-outlet fixed effects (estimates of equation (7)) for sugary (subject to reform) and sugar-free drinks (not subject to reform) and a comparison group consisting of salty licorice, fruit candy, and savory pastry (not subject to the reform). The figure shows that log prices of sugary soda products jump to a higher level of about 0.1 log points right at the reform and relative to sugar-free soda or the control group. Mechanical full pass-through in dashed black line shows how sugary soda prices would have developed with full pass-through of tax to prices, had they otherwise followed the prices of the comparison group post-reform. Vertical line is at the 2014 tax reform. The series are normalized to zero at three weeks prior to the 2014 reform.

Figure 12: Development of quantities, solid products in the 2011 and 2012 reforms

(a) Trends: taxed solids, versus control group and cookies



(b) Event-study: taxed solids vs control, and cookies vs control

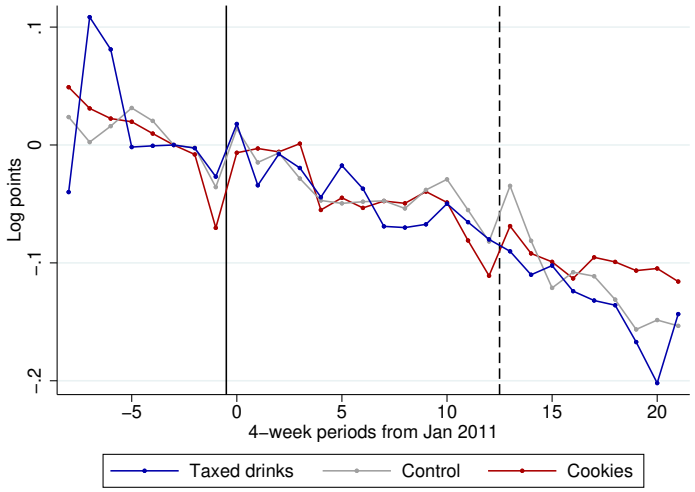


Note: Panel (a) shows regression coefficients for month dummies from regressions of de-seasonalized log number of products sold with product-outlet fixed effects (estimates of equation (7)) for three groups: taxed solids, a control group, and cookies which constitute a substitution group. The three series move similarly around the 2011 and 2012 reforms. Taxable solids include all solid taxable products in the 2011 and 2012 reforms: mainly sweets, chocolate and ice cream. The control group include solid products that were not included in the 2011 and 2012 reforms excluding cookies: mainly sweet and savoury pastry, pudding and yogurt. The series' are normalized to zero three months before the 2011 reform.

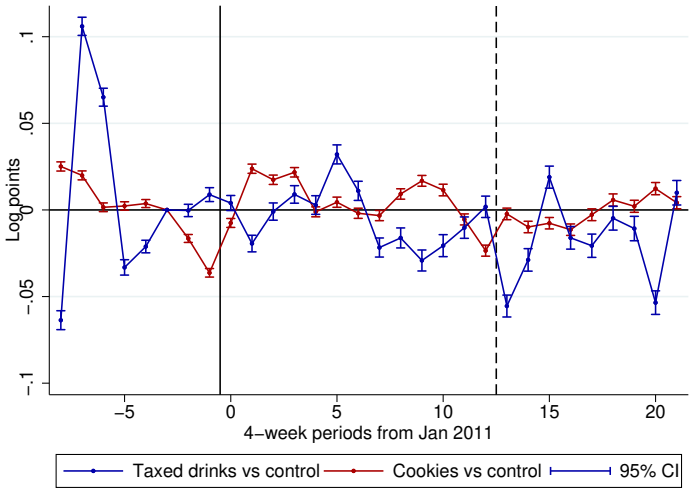
Panel (b) shows the corresponding event-studies comparing taxed solids and cookies to the control group. The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level.

Figure 13: Development of quantities, soft drinks in the 2011 and 2012 reforms

(a) Trends: taxed drinks, versus control group and cookies



(b) Event-study: taxed drinks vs control, and cookies vs control

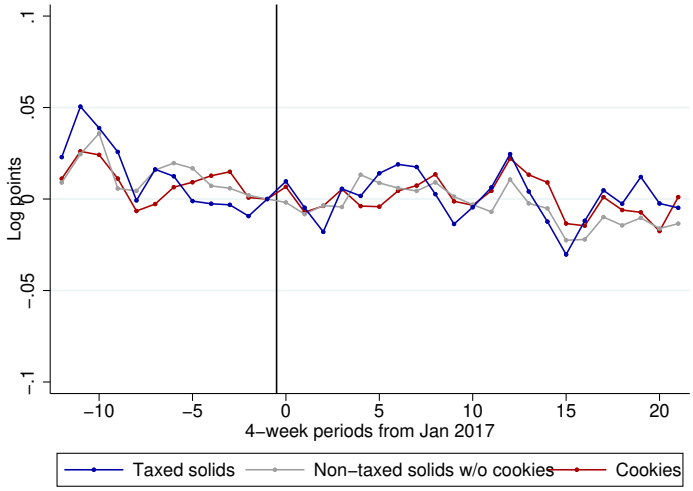


Note: Panel (a) shows regression coefficients for month dummies from regressions of de-seasonalized log number of products sold with product-outlet fixed effects (estimates of equation (7)) for three groups: taxed drinks, a control group, and cookies which constitute a substitution group. The three series move similarly around the 2011 and 2012 reforms. Taxable drinks include all taxable drinks in the 2011 and 2012 reforms: mainly sugary and sugar-free soda and juices. The control group include puddings which were not included in the 2011 and 2012 reforms. The series' are normalized to zero three months before the 2011 reform.

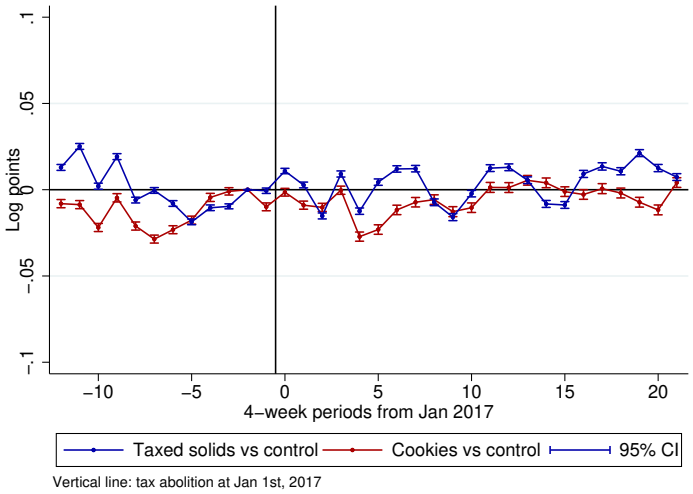
Panel (b) shows the corresponding event-studies comparing taxed drinks and cookies to control groups. The control group for taxed drinks are puddings, while the control group for cookies is all non-taxed solids (excluding cookies). The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level.

Figure 14: Development of quantities, solid products in the 2017 reform

(a) Trends: taxed solids, versus control group and cookies



(b) Event-study: taxed solids vs control, and cookies vs control

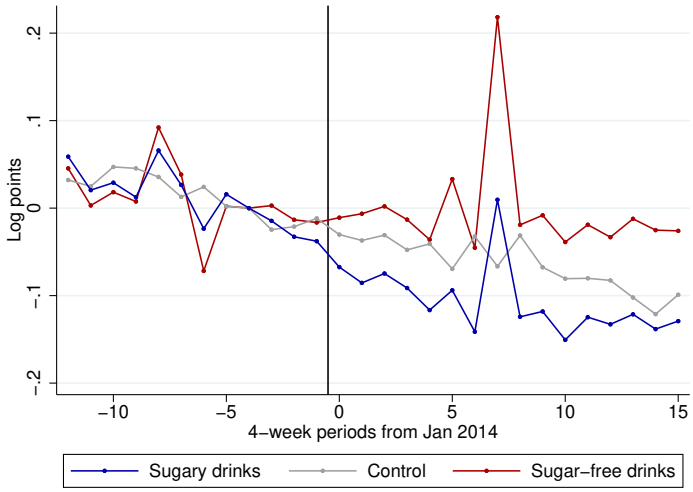


Note: Panel (a) shows regression coefficients for month dummies from regressions of de-seasonalized log number of products sold with product-outlet fixed effects (estimates of equation (7)) for three groups: taxed solids, a control group, and cookies which constitute a substitution group. The three series move similarly around the 2017 reform with some idiosyncratic variation across months. Taxable solids include all solid taxable products in the 2017 reform: mainly sweets, chocolate and ice cream. The control group include solid products that were not included in the 2017 reforms excluding cookies: mainly sweet and savoury pastry, pudding and yogurt. The series' are normalized to zero one month before the 2017 reform.

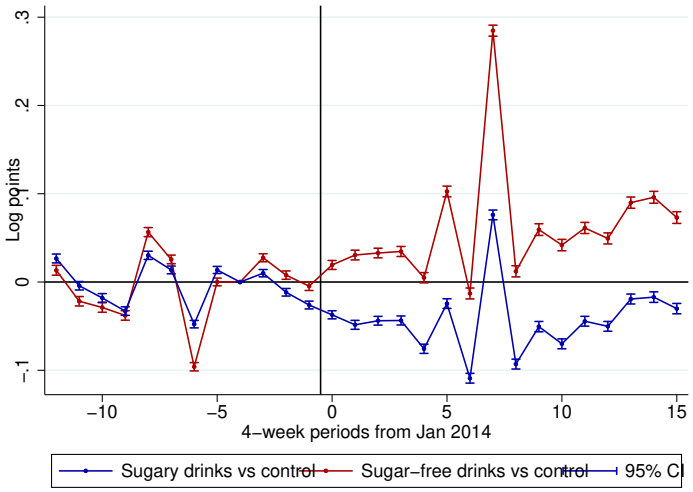
Panel (b) shows the corresponding event-studies comparing taxed solids and cookies to the control group. The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level.

Figure 15: Development of quantities, all sugary and sugar-free drinks in the 2014 reform

(a) Trends: sugary drinks, control group, and sugar-free drinks



(b) Event study: sugary drinks vs control, and sugar-free drinks vs control



Note: Panel (a) shows regression coefficients for month dummies from regressions of de-seasonalized log volume of products sold with product-outlet fixed effects (estimates of equation (7)) for sugary (subject to reform) and sugar-free drinks (not subject to reform) and a control group of honey and salty licorice (not subject to reform). The figure shows that the three series develop in parallel prior to the reform and deviate from each other at the time of the reform such that the sales of sugary drinks decline and sugar-free drinks increase relative to the control group. Vertical line is at the 2014 reform. The series are normalized to zero at four months prior to 2014 reform. The control group consists of honey and salty-licorice, which did not face a tax change in 2014.

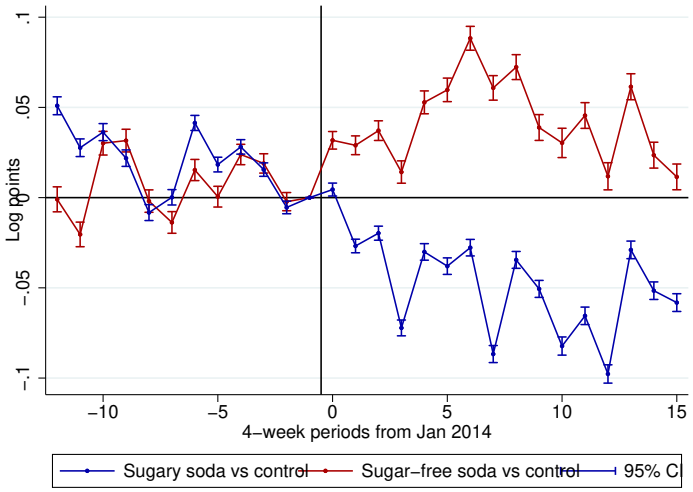
Panel (b) shows the corresponding event-studies comparing sugary drinks and sugar-free drinks to the control group. The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level.

Figure 16: Development of quantities, sugary and sugar-free sodas in the 2014 reform

(a) Trends: sugary soda, control group, and sugar-free soda



(b) Event study: sugary soda vs control, and sugar-free soda vs control

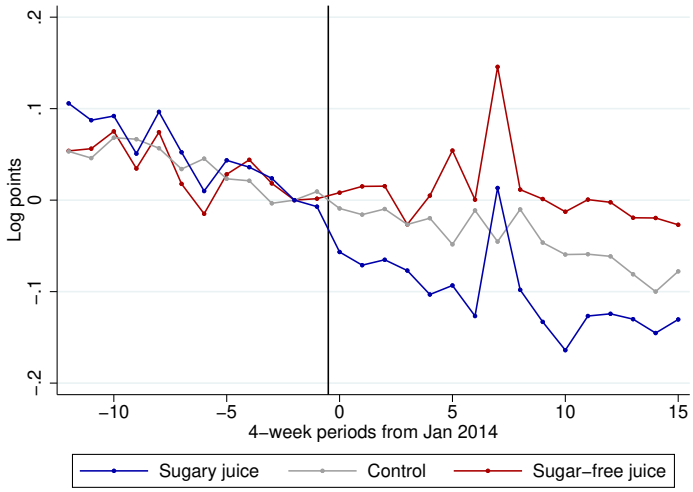


Note: Panel (a) shows regression coefficients for month dummies from regressions of de-seasonalized log volume of products sold with product-outlet fixed effects (estimates of equation (7)) for sugary (subject to reform) and sugar-free soda (not subject to reform) and a control group of ice-cream (not subject to reform). The figure shows that the three series develop in parallel prior to the reform and deviate from each other at the time of the reform such that the sales of sugary soda decline and sugar-free soda increase relative to the control group. Vertical line is at the 2014 reform. The series are normalized to zero at two months prior to 2014 reform. The control group consists of ice-creams, which did not face a tax change in 2014.

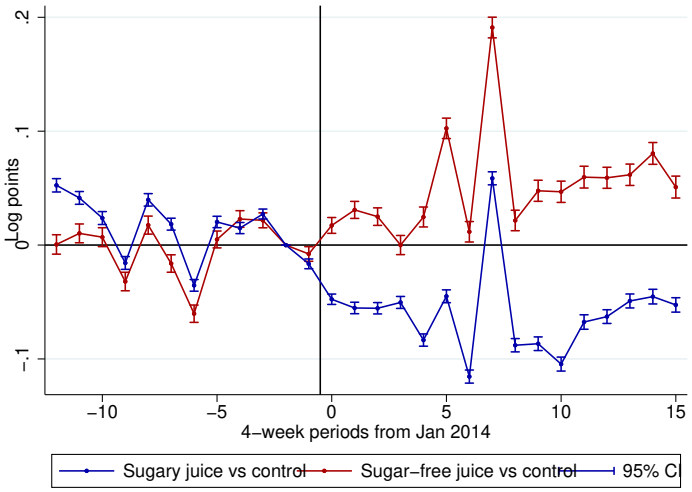
Panel (b) shows the corresponding event-studies comparing sugary sodas and sugar-free sodas to the control group. The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level.

Figure 17: Development of quantities, sugary and sugar-free juices in the 2014 reform

(a) Trends: sugary juice, control group, and sugar-free juice



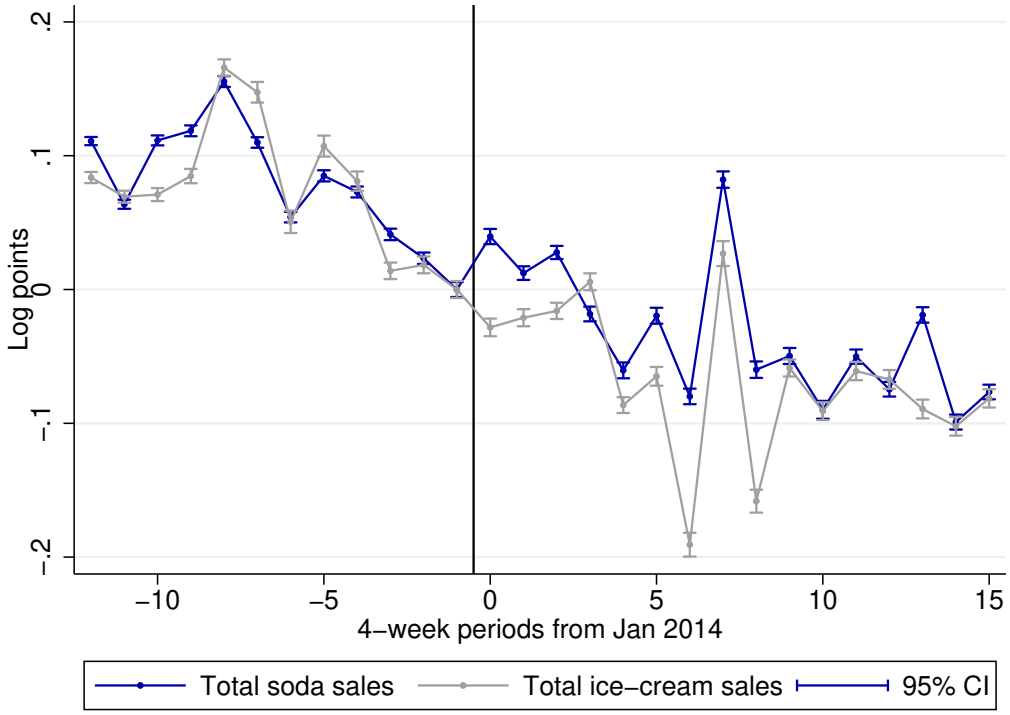
(b) Event study: sugary juice vs control, and sugar-free juice vs control



Note: Panel (a) shows regression coefficients for month dummies from regressions of de-seasonalized log volume of products sold with product-outlet fixed effects (estimates of equation (7)) for sugary (subject to reform) and sugar-free juices and juice concentrates (not subject to reform) and a control group of honey and salty licorice (not subject to reform). The figure shows that the three series develop in parallel prior to the reform and deviate from each other at the time of the reform such that the sales of sugary juice decline and sugar-free juice increase relative to the control group. Vertical line is at the 2014 reform. The series are normalized to zero at two months prior to 2014 reform. The control group consists of ice-creams, which did not face a tax change in 2014.

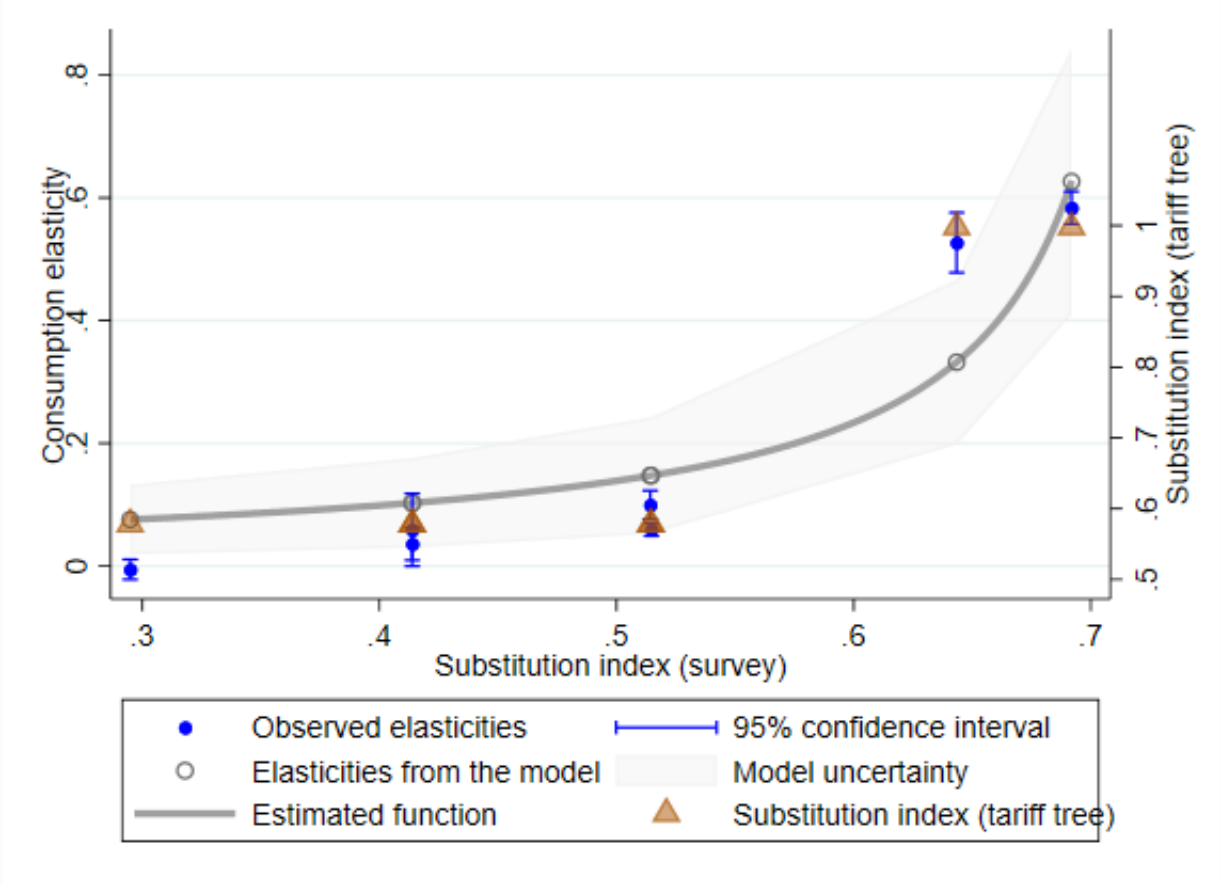
Panel (b) shows the corresponding event-studies comparing sugary juices and sugar-free juices to the control group. The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level.

Figure 18: Development of total soda sales in the 2014 reform



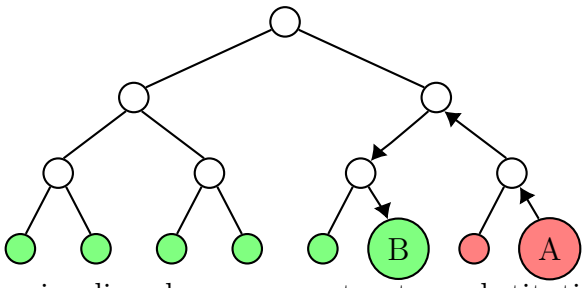
Note: The figure shows regression coefficients for month dummies from regressions of log volume of products sold with product-outlet fixed effects (estimates of equation (7)). The figure shows that the two groups develop in parallel both before and after the reform such that the sales of soda sales in total do not decline after the reform relative to the control group. The observations have been weighted by total volumes sold of each product at each outlet in 2013, so that the estimates approximately correspond to the development of log total volumes. The figure includes 95% confidence intervals surrounding the coefficients based on standard errors that are clustered at the product-outlet level. Vertical line is at the 2014 reform. The series are normalized to zero at five months prior to 2014 reform. The control group consists of ice-creams, which did not face a tax change in 2014.

Figure 19: Fitted values of the model



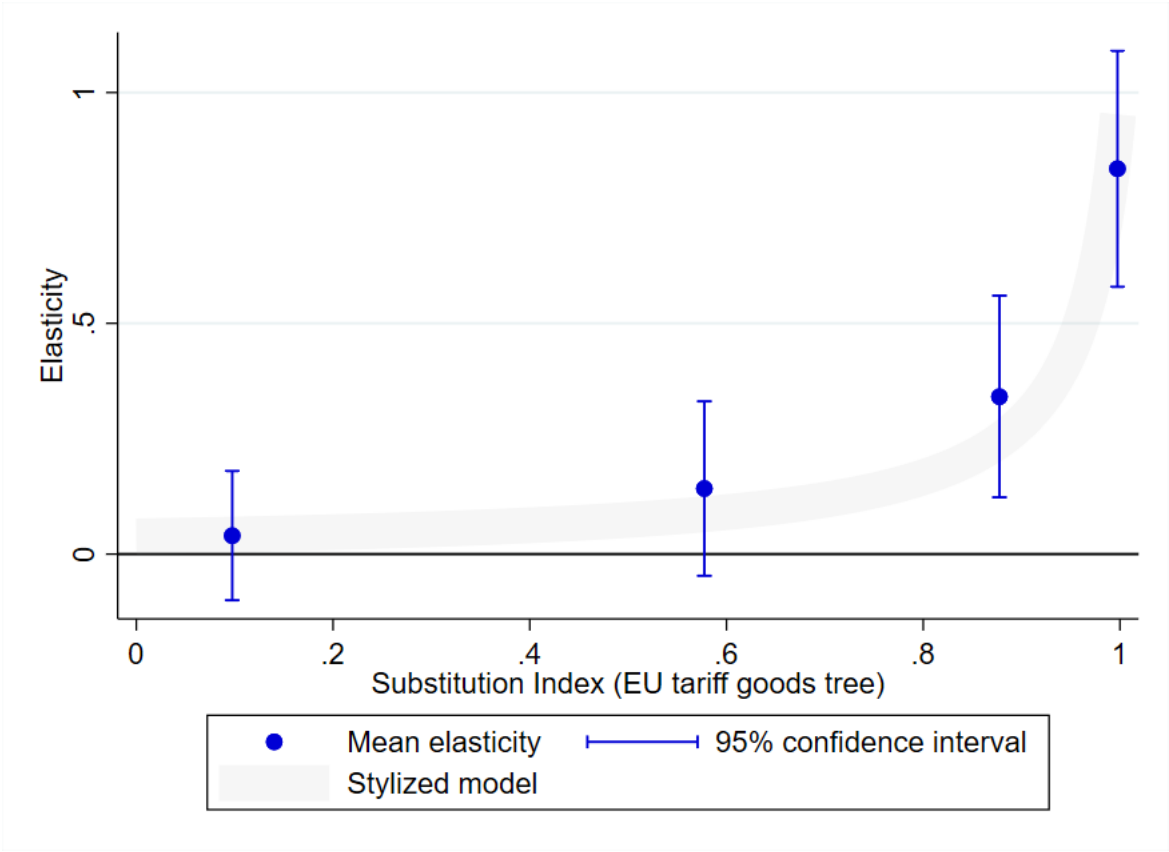
Note: The figure shows estimates of consumption elasticities for different subgroups of taxed products from Table 3 as well as estimates for sugary and sugar-free drinks in the 2011 and 2012 reforms (pooled). The figure also shows fitted values from an empirically applicable version of our model organized by the substitution index of the product group facing tax changes and the closest product-group not facing tax changes in the horizontal axis. Substitution indices based on the EU tariff tree are also shown (right axis). The figure shows that the consumption elasticity grows non-linearly and becomes meaningfully higher than zero only when the substitution index reaches a level of about 0.6 with the highest level being 1. This convex pattern holds for both the estimated elasticities and the prediction from the theory model.

Figure 20: Illustration of a substitution index based on the EU-tariff tree



Notes: This graph visualizes how we construct a substitution index from the EU-tariff goods tree. Each level in the tree here represents a two-digit category level, with the highest node corresponding to a category encompassing all goods in the tree. Here, we note taxed categories by red and non-taxed categories by green. As an example, when we construct the substitution index for goods in the category A, then B is a closest non-taxed category. The shortest route from A to B goes through two levels of categories. The longest possible route goes through three levels, the length of which we normalize to 1. Hence, the normalized distance (d) from A to B is 2/3. We define the substitution index as $1 - d^2$, where the square comes from the idea that the number of plausible substitution categories at a similar distance increases approximately to the power of two.

Figure 21: Elasticities from the literature versus the substitution index from the EU tariff goods tree



Note: This figure plots meta-analysis estimates of elasticities from previous literature in four categories using a DerSimonian-Laird random-effects model. The figure shows a highly convex pattern similar to our estimates from the sweets tax scheme and aligns well witht the prediction from our theory model presented in a shaded grey curve underneath the estimated elasticities from the literature. The categories are based on the distance to a closest substitute with no price changes based on the EU tariff goods tree. This distance is then formed into a substitution index with values between 0 (not substitutable at all) and 1 (complete substitute, i.e. same good) (see section 6 for details on how the substitution index is formed).

Tables

Table 1: Descriptive Statistics

	Taxed			Non-taxed		
	Mean	SD	N	Mean	SD	N
Panel A: Full sample, all taxed vs. non-taxed products						
Price (€)	1.92	1.33	330,493,591	2.21	2.01	213,915,121
Quantity (units)	11.84	28.92	330,493,591	15.17	73.02	213,915,121
Panel B: Pre-2011/2012 reforms (2010), all taxed vs. non-taxed products						
Price (€)	1.56	1.09	33,104,327	2.19	2.03	21,075,956
Quantity (units)	13.28	33.64	33,104,327	14.63	36.66	21,075,956
Panel C: Pre-2014 differentiation (2013), taxed vs. non-taxed drinks						
Price (€)	1.84	1.12	11,828,885	1.80	1.11	4,281,982
Quantity (units)	12.15	27.03	11,828,885	12.68	32.54	4,281,982
Panel D: Pre-2017 abolition (2016), taxed vs. non-taxed solids						
Price (€)	1.92	1.34	20,189,664	2.10	1.77	25,689,193
Quantity (units)	11.45	30.97	20,189,664	16.54	99.82	25,689,193

Notes: This table reports descriptive statistics for scanner data from S-group stores. Prices are in euros per unit. Quantity refers to weekly units sold per product-store observation. All taxed products include e.g. sweets, chocolate, ice cream, and soft drinks. All non-taxed products include e.g. cookies, pastries, pudding, and yogurt. Panel A shows descriptive statistics for the full sample from the beginning of 2010 to the end of September 2018. Panel B shows descriptive statistics for all taxed products (both solid and drinks) and non-taxed products before the 2011/2012 reforms in 2010. Panel C shows descriptive statistics for taxed and non-taxed drinks before the 2014 reform that separated the taxation of sugary and sugar-free drinks. Panel D shows descriptive statistics for taxed and non-taxed solids before the 2017 abolition of the tax on solid products.

Table 2: Price and quantity changes of reformed and control groups in different reforms

	2011 & 2012 reforms		**2017 reform**		**2014 reform**	
	Treatment: Taxed solids	Substitution: Cookies	Treatment: Taxed solids	Substitution: Cookies	Treatment: Sugary liquids	Substitution: Sugar-free liquids
	(1)	(2)	(3)	(4)	(5)	(6)
Price change	0.150	0.026	-0.093	0.003	0.088	-0.016
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Quantity change	-0.002	0.002	-0.002	0.006	-0.041	0.065
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Control group	Non-taxed solids w/o cookies	Non-taxed solids w/o cookies	Non-taxed solids w/o cookies	Non-taxed solids w/o cookies	Honey & salty licorice	Honey & salty licorice
N	23,186,261	7,496,341	21,418,727	13,318,485	7,852,502	3,440,104
4-week FE's	X	X	X	X	X	X
Product–outlet FE's	X	X	X	X	X	X

Notes: This table provides difference-in-differences estimates of the quantity and price responses and associated elasticities of taxed solids in the 2011-2012 reforms. Estimate for quantities is -0.004, estimate for prices is 0.143, and estimate for elasticity is -0.028. The control group is non-taxed solids without cookies. The estimates compare seasonalized log outcomes pre-reform (2010) to post-reform (2012). There are two specifications: (1) controls for 4-week fixed effects, while (2) adds product–outlet fixed effects. Outcomes are seasonalized, and standard errors are clustered at the product–outlet level.

Table 3: Quantity and price changes and elasticities by sub-groups, preferred estimates

	Sweets	Chocolate	Ice-cream	Sugary soda	Sugary juice
Quantity change	0.002	-0.007	-0.005	-0.040	-0.079
	(0.000)	(0.000)	(0.000)	(0.002)	(0.002)
Price change	0.089	0.088	0.159	0.075	0.135
	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)
Elasticity	0.022	-0.080	-0.031	-0.533	-0.584
	(0.004)	(0.005)	(0.002)	(0.036)	(0.013)
Substitution index	0.295	0.514	0.515	0.643	0.692
(survey)	(0.006)	(0.006)	(0.018)	(0.006)	(0.008)
Substitution index	0.578	0.578	0.578	0.998	0.998
(tariff tree)					
N (quantity, price,	24,346,341	19,833,349	21,221,339	3,449,246	2,724,312
and elasticity)					
N (Survey)	2,245	2,241	528	2,830	1,525
Control group	Non-taxed solids	Non-taxed solids	Non-taxed solids	Ice-cream	Honey &
	w/o cookies	w/o cookies	w/o cookies		salty licorice
Tax reform	2011, 2012	2011, 2012	2011, 2012	2014	2014
	& 2017	& 2017	& 2017		

Notes: This table provides our preferred difference-in-differences estimates of the quantity and price responses and associated elasticities of taxed product categories. The sub-groups are: sweets, ice-cream, chocolate, sugary soda, and sugary juice. Elasticity estimates are 0.022, -0.031, -0.080, -0.533, and -0.584 respectively. The estimates compare de-seasonalized log outcomes one year before a given reform to one year post-reform. Pooled estimates for the solids (sweets, ice-cream and chocolate) compare all reforms such that the 2011 reform is normalized to be 0.75 €/kg / 0.95 €/kg = 0.789 of the full reform. Standard errors are clustered at the product–outlet level.

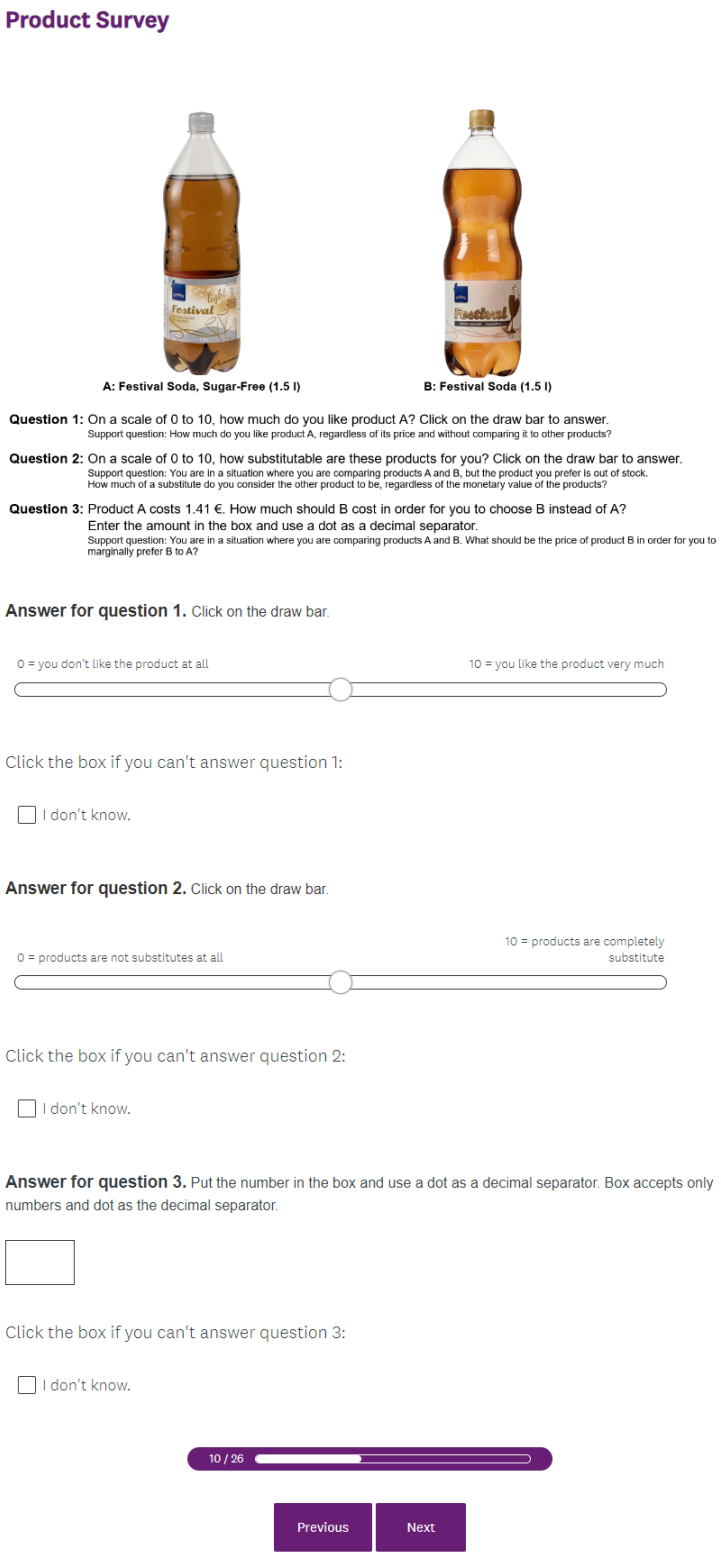
Table 4: Studies included in the meta-analysis

Study	Good	Closest substitute	Cat.	SI	Elasticity
Allcott and Rafkin (2022)	Cigarettes	Other tobacco prod.	0	0.9975	1.908
Allcott and Rafkin (2022)	E-cigarettes	Other nicotine prod.	0	0.9975	1.318
Cawley et al. (2020)	Sweetened beverages	Non-sweetened beverages	0*	0.9975	1.300
Chiou and Muehlegger (2014)	Cigarettes	Other tobacco prod.	0	0.9975	0.355
Colantuoni and Rojas (2015)	Sweetened beverages	Non-sweetened beverages	0*	0.9975	0.250
Colantuoni and Rojas (2015)	Sweetened beverages	Non-sweetened beverages	0*	0.9975	0.154
Cotti et al. (2016)	Chewing tobacco	Other tobacco prod.	0	0.9975	0.206
Cotti et al. (2016)	Snuff	Other tobacco prod.	0	0.9975	0.849
Cotti et al. (2016)	Cigarettes	Other tobacco prod.	0	0.9975	2.200
Cotti et al. (2016)	E-cigarettes	Other nicotine prod.	0	0.9975	0.440
Evans et al. (1999)	Cigarettes	Other tobacco prod.	0	0.9975	0.294
Gruber et al. (2003)	Cigarettes	Other tobacco prod.	0	0.9975	0.450
Harding et al. (2012)	Cigarettes	Other tobacco prod.	0	0.9975	0.369
Lovenheim (2008)	Cigarettes	Other tobacco prod.	0	0.9975	0.457
Powell et al. (2020)	Sweetened beverages	Non-sweetened beverages	0*	0.9975	1.100
Roberto et al. (2019)	Sweetened beverages	Non-sweetened beverages	0*	0.9975	1.700
Rojas and Wang (2021)	Carbonated drinks	Non-carbonated drinks	0*	0.9975	0.970
Seiler et al. (2021)	Sweetened beverages	Non-sweetened beverages	0*	0.9975	0.650
Asplund et al. (2007)	Beer	Non-alcoholic drinks	1*	0.8775	0.467
Asplund et al. (2007)	Spirits	Non-alcoholic drinks	1*	0.8775	0.321
Asplund et al. (2007)	Wine	Non-alcoholic drinks	1*	0.8775	1.720
Gonçalves and dos Santos (2020)	Sweetened beverages	Non-sweetened beverages	1	0.8775	0.570
Gillingham et al. (2015)	Gasoline	Other forms of energy	2	0.5775	0.099
Hindriks and Serse (2022)	Electricity	Other forms of energy	2	0.5775	0.017
Langer et al. (2017)	Mile of travel by car	Other forms of transport	2	0.5775	0.150
Levin et al. (2017)	Gasoline	Other forms of energy	2	0.5775	0.301
Harju and Kosonen (2010)	Restaurant meals	Home-made meals	3	0.0975	-0.045
Kosonen (2015)	Haircuts	Home haircuts	3	0.0975	0.082
Merker (2023)	Sweets and beverages	Other food	3	0.0975	0.083

Note: Cat. refers to number of categories between taxed good and closest substitute good in the EU tariff tree. Category values with an asterisk (*) are reduced by one due to substantial possibilities for cross-border shopping. SI refers to a EU tariff goods tree based substitution index.

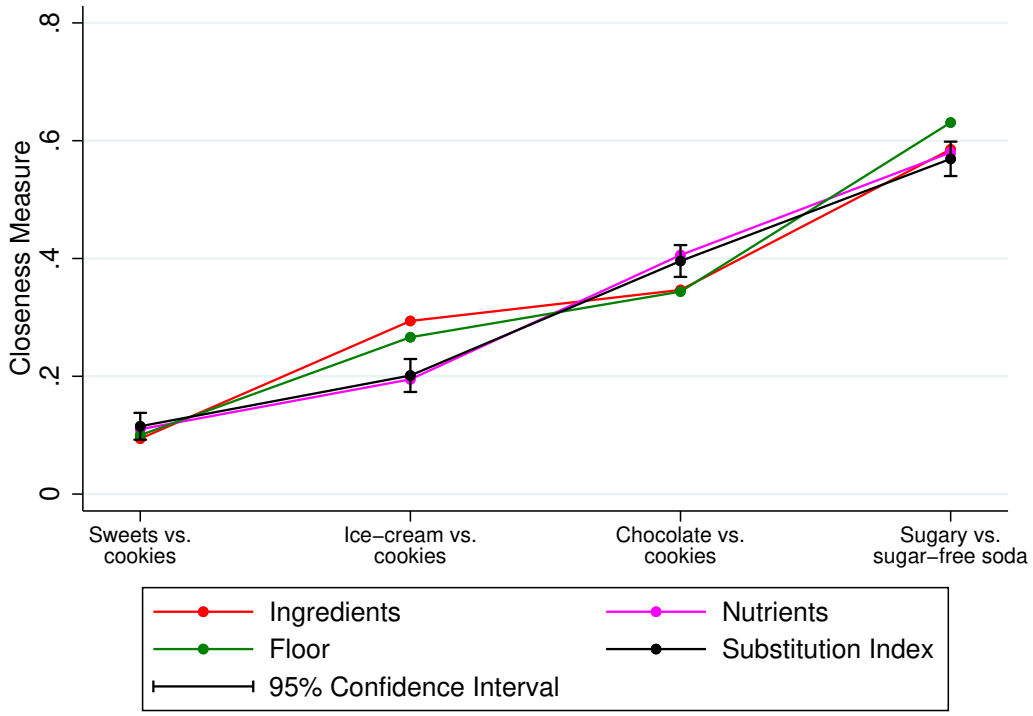
A APPENDIX - Example of the substitution survey

Figure A.1: Example of the substitution survey



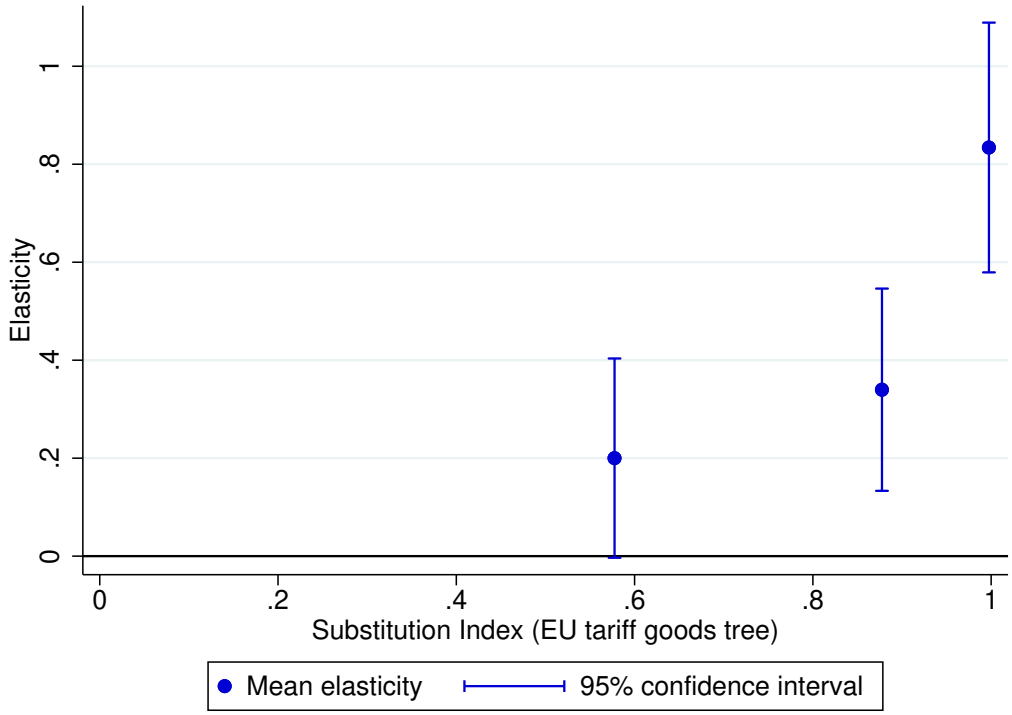
B APPENDIX - Additional figures and tables

Figure A.2: Alternative measures of substitutability between taxed and their closest non-taxed goods category



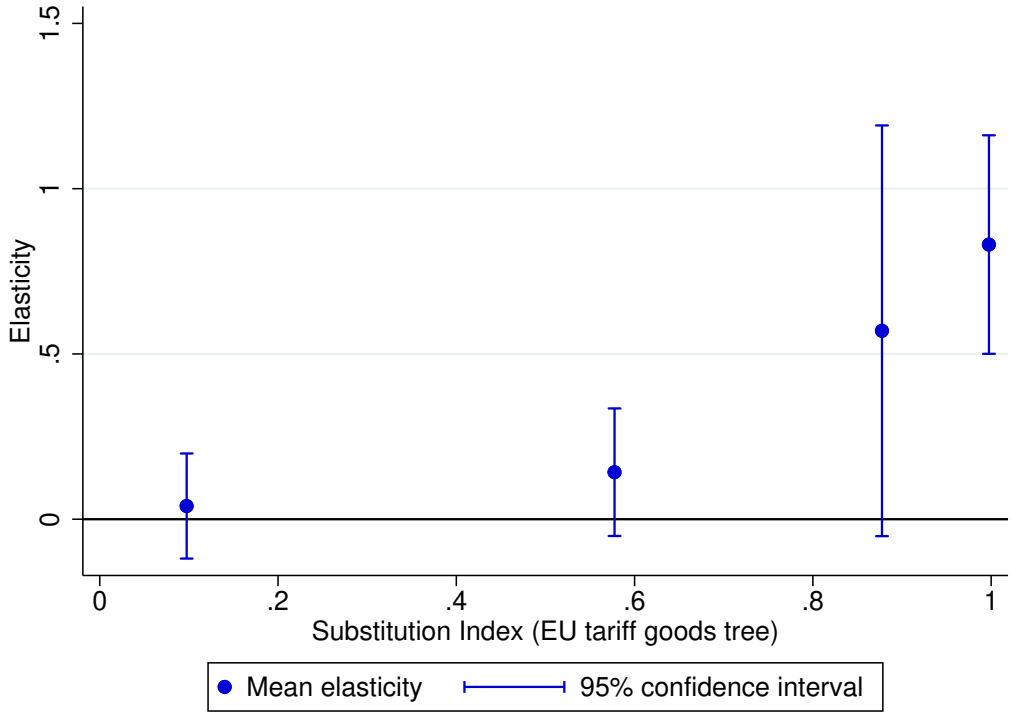
Note: The figure shows mean values of four similarity measures: the survey based substitution index and measures related to the ingredients, nutrient content and store floor location of products. The similarity measures are relative to the closest low-tax category for each product group. The closest low-tax category is chosen by the substitution index and is cookies for sweets, ice-cream and chocolate, and sugar-free soda for sugary soda. The higher the value of a measure, the more closer the categories are to their low-tax counterpart. 95% confidence intervals are also included for the substitution index.

Figure A.3: Elasticities from the literature versus the substitution index from the EU tariff goods tree. Robustness: Excluding estimates for services and unpublished research.



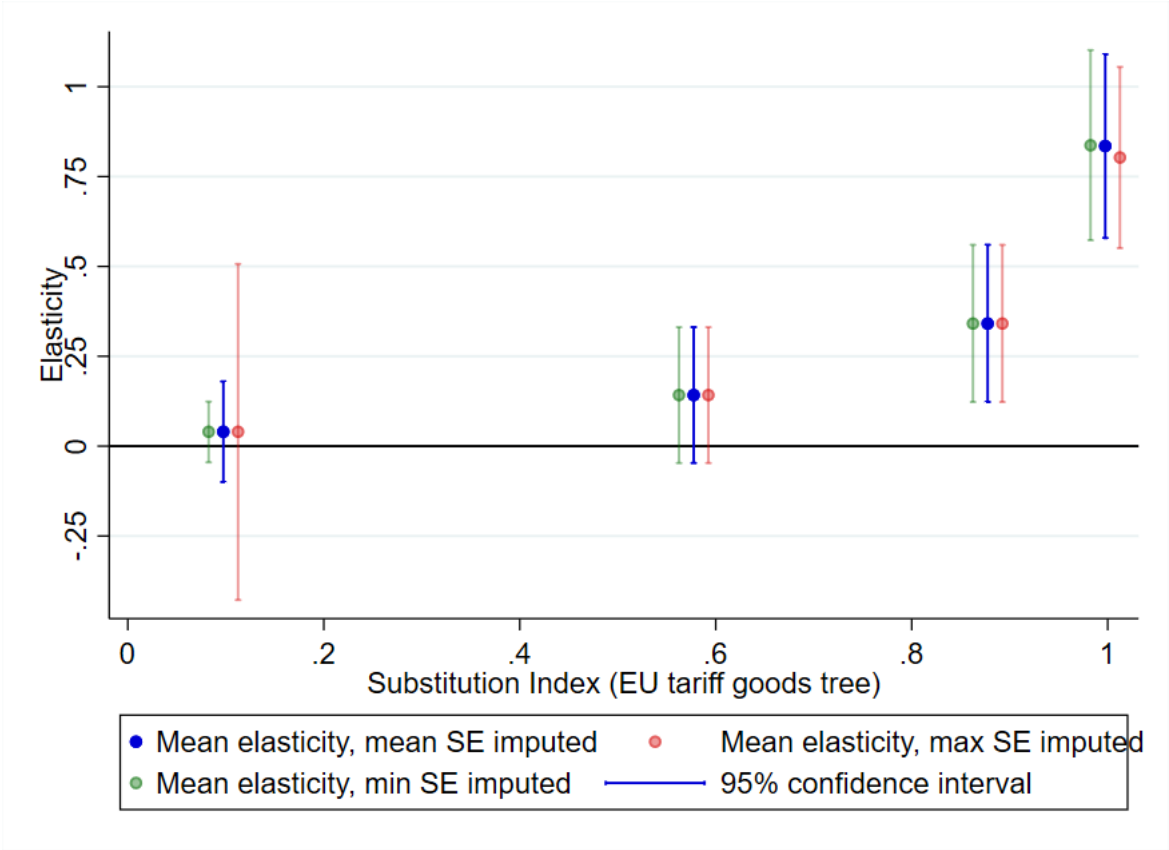
Note: This figure plots meta-analysis results for elasticities from previous literature as in Figure 21, but excluding estimates for services and unpublished work. Only three out of the four categories have studies with this exclusion. The category-level estimates are from the DerSimonian-Laird random-effects model. The categories are based on the distance to a closest substitute with no price changes based on the EU tariff goods tree. This distance is then formed into a substitution index with values between 0 (not substitutable at all) and 1 (complete substitute, i.e. same good) (see section 6 for details on how the substitution index is formed).

Figure A.4: Elasticities from the literature versus the substitution index from the EU tariff goods tree. Robustness: Excluding estimates with adjustment for cross-border shopping.



Note: This figure plots meta-analysis results for elasticities from previous literature as in Figure 21, but excluding estimates for which the substitution index was adjusted for cross-border shopping. The category-level estimates are from the DerSimonian-Laird random-effects model. The categories are based on the distance to a closest substitute with no price changes based on the EU tariff goods tree. This distance is then formed into a substitution index with values between 0 (not substitutable at all) and 1 (complete substitute, i.e. same good) (see section 6 for details on how the substitution index is formed).

Figure A.5: Elasticities from the literature versus the substitution index from the EU tariff goods tree. Robustness: SE imputation for missing SE's



Note: This figure plots meta-analysis results for elasticities from previous literature in four categories using a DerSimonian-Laird random-effects model. The categories are based on the distance to a closest substitute with no price changes based on the EU tariff goods tree. This distance is then formed into a substitution index with values between 0 (not substitutable at all) and 1 (complete substitute, i.e. same good) (see section 6 for details on how the substitution index is formed). Different series represent different meta-analysis estimates of elasticities for the groups based on different assumptions made on missing standard error values. These are either imputed to match the minimum, mean or maximum standardized standard errors of the sample of estimates that report standard errors.

C APPENDIX - Derivations for the general model

Consider a Gorman-Lancaster model (Gorman 1980, Lancaster 1966) where an agent derives utility from the characteristics of the goods she consumes, not from the goods themselves. Assume that there are n goods and $m \leq n$ characteristics. If there are more characteristics then goods, the characteristics representation is not feasible (Lancaster 1966). Moreover, note that the consumer will consume only m goods in equilibrium, so that we will proceed as if there were only m goods.

The utility function of the agent is then given by $u(c_1, c_2, ..., c_m) = u(c)$, where c is a vector of characteristics consumption. The vector is formed by consuming goods $x = (x_1, x_2, ..., x_m)'$ so that:

$$c = Cx,$$

where $C = \begin{bmatrix} c_{11} & c_{12} & \cdots & c_{1m} \\ c_{21} & c_{22} & \cdots & c_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ c_{m1} & c_{m2} & \cdots & c_{mm} \end{bmatrix}$ is a $m \times m$ matrix, and c_{ij} is the characteristics coefficient of good j to characteristic i. Assuming that C is invertible[16], we have:

$$x = C^{-1}c. \tag{12}$$

As there is an equal number of goods and characteristics, there exists a one-to-one mapping between the characteristics space and the goods space. Hence, the problem of the consumer can be written as:

$$\max_c u(c) \, s.t. \, p'x = w,$$

where $p = (p_1, p_2, ..., p_m)'$ is a vector of product prices, and w is the level of income.

[16]This is equivalent to assuming that none of the characteristics vectors of the goods is a linear combination of the others, so that none of the goods are redundant.

Role of closeness

Consider a situation where good 2 moves closer to good 1 by moving on the budget surface in the characteristics space. This is achieved if the characteristics vector of good 2 becomes $(\alpha r c_{11} + (1-\alpha)c_{12}, ..., \alpha r c_{m1} + (1-\alpha)c_{m2})$ for some $\alpha \in [0, 1[$, and where $r = \frac{p_2}{p_1}$ is the price ratio of goods 2 and 1 before any price changes. We denote the new characteristics matrix in this situation by $C(\alpha)$. Note that $C = C(\alpha|\alpha = 0)$ is the original characteristics matrix.

Expression for $\frac{\partial x(\alpha)}{\partial \alpha}$

Notice also that the gradient of the budget surface in the characteristics space does not change when α changes. Hence, assuming the change does not result in a corner solution, we have $\frac{\partial c}{\partial \alpha} = 0$. As $c = C(\alpha)x(\alpha)$, we have that:

$$\frac{\partial C(\alpha)x(\alpha)}{\partial \alpha} = \frac{\partial C(\alpha)}{\partial \alpha}x(\alpha) + C(\alpha)\frac{\partial x(\alpha)}{\partial \alpha} = 0.$$

Re-arranging and pre-multiplying by $C(\alpha)^{-1}$, we have:

$$\frac{\partial x(\alpha)}{\partial \alpha} = -C(\alpha)^{-1}\frac{\partial C(\alpha)}{\partial \alpha}x(\alpha). \tag{13}$$

Note here that $\frac{\partial C(\alpha)}{\partial \alpha} = \begin{bmatrix} 0 & rc_{11} - c_{12} & 0 & \cdots & 0 \\ 0 & rc_{21} - c_{22} & 0 & \cdots & 0 \\ 0 & rc_{31} - c_{32} & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & rc_{m1} - c_{m2} & 0 & \cdots & 0 \end{bmatrix} = r[0_{m\times 1}, C_{(1)}(\alpha), 0_{m\times 1}, ..., 0_{m\times 1}] -$
$[0_{m\times 1}, C_{(2)}, 0_{m\times 1}, ..., 0_{m\times 1}]$ where $C_{(1)}(\alpha)$ is the first column of $C(\alpha)$ and $C_{(2)}$ is the second column of $C = C(\alpha|\alpha = 0)$. Hence, we have that:

$$\begin{aligned} C(\alpha)^{-1}\frac{\partial C(\alpha)}{\partial \alpha} = r[&C(\alpha)^{-1}0_{m\times 1}, C(\alpha)^{-1}C_{(1)}(\alpha), C(\alpha)^{-1}0_{m\times 1}, ..., C(\alpha)^{-1}0_{m\times 1}] \\ &- [C(\alpha)^{-1}0_{m\times 1}, C(\alpha)^{-1}C_{(2)}, C(\alpha)^{-1}0_{m\times 1}, ..., C(\alpha)^{-1}0_{m\times 1}]. \end{aligned}$$

So that, as $C(\alpha)^{-1}0_{m\times 1} = 0_{m\times 1}$ and $C(\alpha)^{-1}C_{(1)}(\alpha) = (1, 0, ..., 0)'$, we have:

$$C(\alpha)^{-1}\frac{\partial C(\alpha)}{\partial \alpha} = \begin{bmatrix} 0 & r & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & 0 \end{bmatrix} - [0_{m\times 1}, C(\alpha)^{-1}C_{(2)}, 0_{m\times 1}, ..., 0_{m\times 1}]. \tag{14}$$

Hence, we are left with the problem of finding a closed form expression for $C(\alpha)^{-1}C_{(2)}$. Notice that the second column of $C(\alpha)$ is $C_{(2)}(\alpha) = \alpha r C_{(1)}(\alpha) + (1-\alpha)C_{(2)}$, so that:

$$C(\alpha)^{-1}C_{(2)}(\alpha) = \alpha r C(\alpha)^{-1}C_{(1)}(\alpha) + (1-\alpha)C(\alpha)^{-1}C_{(2)}.$$

Re-arranging, we have that:

$$C(\alpha)^{-1}C_{(2)} = \frac{1}{1-\alpha}C(\alpha)^{-1}C_{(2)}(\alpha) - \frac{\alpha}{1-\alpha}rC(\alpha)^{-1}C_{(1)}(\alpha).$$

As $C(\alpha)^{-1}C_{(2)}(\alpha) = [0, 1, 0, ..., 0]'$ and $C(\alpha)^{-1}C_{(1)}(\alpha) = [1, 0, ..., 0]'$, we get:

$$[0_{m\times 1}, C(\alpha)^{-1}C_{(2)}, 0_{m\times 1}, ..., 0_{m\times 1}] = \begin{bmatrix} 0 & -\frac{\alpha}{1-\alpha}r & \cdots & 0 \\ 0 & \frac{1}{1-\alpha} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & 0 \end{bmatrix}.$$

Inserting this to 14, we have that:

$$C(\alpha)^{-1}\frac{\partial C(\alpha)}{\partial \alpha} = \begin{bmatrix} 0 & \frac{1}{1-\alpha}r & \cdots & 0 \\ 0 & -\frac{1}{1-\alpha} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & 0 \end{bmatrix}.$$

So that inserting this expression to 13 we have:

$$\frac{\partial x(\alpha)}{\partial \alpha} = \begin{bmatrix} 0 & -\frac{1}{1-\alpha}r & \cdots & 0 \\ 0 & \frac{1}{1-\alpha} & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & 0 \end{bmatrix} x(\alpha) = \begin{bmatrix} -\frac{1}{1-\alpha}rx_2(\alpha) \\ \frac{1}{1-\alpha}x_2(\alpha) \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \tag{15}$$

Note that, by 12, $x_2(\alpha) = C(\alpha)_2^{-1}c$ where $C(\alpha)_2^{-1}$ is the second row of $C(\alpha)^{-1}$. Using Cramer's rule, we have that:

$$x_2(\alpha) = C(\alpha)_2^{-1}c = \frac{1}{detC(\alpha)}[D_{12}(\alpha), D_{22}(\alpha), ..., D_{m2}(\alpha)]c, \tag{16}$$

where $D_{ij}(\alpha)$ is the (i,j)-th entry of the cofactor matrix of $C(\alpha)$. As a cofactor $D_{ij}(\alpha)$ depends only on the entries outside the i-th row and j-th column of $C(\alpha)$, none of the cofactors in 16 depend on the second column of $C(\alpha)$. As only the second column of $C(\alpha)$ depends on α, none of the cofactors in 16 depend on α. I.e. $D_{i2}(\alpha) = D_{i2}(\alpha|\alpha = 0) = D_{i2}$. Moreover, note that $C(\alpha)$ can be constructed from C by two elementary column operations: 1) multiplying the second column of C by $(1-\alpha)$, and 2) adding the first column multiplied by αr to the second column. As the first operation leads to the multiplication of the determinant by $(1-\alpha)$ and the second operation does not change the determinant, $detC(\alpha) = (1-\alpha)detC$. Therefore, we have from 16 that:

$$\begin{aligned} x_2(\alpha) &= C(\alpha)_2^{-1}c \\ &= \frac{1}{detC(\alpha)}[D_{12}(\alpha), D_{22}(\alpha), ..., D_{m2}(\alpha)]c \\ &= \frac{1}{1-\alpha}\frac{1}{detC}[D_{12}, D_{22}, ..., D_{m2}]c \\ &= \frac{1}{1-\alpha}C(\alpha|\alpha = 0)_2^{-1}c \\ &= \frac{1}{1-\alpha}x_2(\alpha|\alpha = 0) \\ &= \frac{1}{1-\alpha}x_2. \end{aligned} \tag{17}$$

Combining 17 with 15, we arrive at an expression for $\frac{\partial x}{\partial \alpha}$ based solely on α, the original price ratio r, and the demand for good 2 when $\alpha = 0$:

$$\frac{\partial x(\alpha)}{\partial \alpha} = \begin{bmatrix} -\frac{1}{(1-\alpha)^2} r x_2 \\ \frac{1}{(1-\alpha)^2} x_2 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \tag{18}$$

Expression for $x(\alpha)$

Note that, by 18, the demands for all the other goods apart from 1 and 2 are constant with respect to α. Hence, when α changes, the budget share of those goods remains the same. This implies that the total budget share of goods 1 and 2 stays the same, so that:

$$p_1 x_1 + p_2 x_2 = p_1(x_1 + \int_0^\alpha \frac{\partial x_1(h)}{\partial h} dh) + p_2(x_2 + \int_0^\alpha \frac{\partial x_2(h)}{\partial h} dh).$$

Simplifying and using information from 18, we have:

$$p_1 \int_0^\alpha \frac{\partial x_1(h)}{\partial h} dh = p_2 x_2 \int_0^\alpha \frac{1}{(1-h)^2} dh$$

so that, by integration, we have:

$$p_1(x_1(\alpha) - x_1) = \left(\frac{1}{1-\alpha} - 1\right) p_2 x_2.$$

Re-arranging it gives us:

$$x_1(\alpha) = x_1 - \frac{\alpha}{1-\alpha} r x_2.$$

Combining this information with 17 and, as the demand for other goods is not affected by α, we have:

$$x(\alpha) = \begin{bmatrix} x_1 - \frac{\alpha}{1-\alpha} r x_2 \\ \frac{1}{1-\alpha} x_2 \\ x_3 \\ \vdots \\ x_m \end{bmatrix} \tag{19}$$

Expression for $\varepsilon_1(\alpha) \equiv \frac{\partial x_1(\alpha)}{\partial p_1} \frac{p_1}{x_1(\alpha)}$

From 19 we have that:

$$\frac{\partial x_1(\alpha)}{\partial p_1} = \frac{\partial x_1}{\partial p_1} - \frac{\alpha}{1-\alpha} r \frac{\partial x_2}{\partial p_1}$$

as r here is the original price ratio (and thus does not change). Multiplying with p_1 and dividing by $x_1(\alpha)$ from equation 19, we have:

$$\varepsilon_1(\alpha) = \frac{\frac{\partial x_1}{\partial p_1} p_1 - \frac{\alpha}{1-\alpha} \frac{\partial x_2}{\partial p_1} p_2}{x_1 - \frac{\alpha}{1-\alpha} \frac{p_2}{p_1} x_2}.$$

Then, dividing both the numerator and the denominator by $(1-\alpha)x_1$ on the RHS and re-arranging, we have:

$$\varepsilon_1(\alpha) = \frac{(1-\alpha)\varepsilon_1 - \alpha \frac{b_2}{b_1} \varepsilon_{21}}{1 - \alpha - \alpha \frac{b_2}{b_1}},$$

where b_i is the budget share of good i. Decomposing ε_1 and ε_{21} by the Slutsky equation, we then have:

$$\varepsilon_1(\alpha) = \frac{\varepsilon_1^h - \alpha(\frac{b_2}{b_1} \varepsilon_{21}^h + \varepsilon_1^h)}{1 - \alpha - \alpha \frac{b_2}{b_1}} - \frac{(1-\alpha)\varepsilon_1^w b_1 - \alpha \varepsilon_2^w b_2}{1 - \alpha - \alpha \frac{b_2}{b_1}}.$$

Noting that, by the budget constraint $\frac{b_2}{b_1} \varepsilon_{21}^h + \varepsilon_1^h = -\sum_{i=3}^m \frac{b_i}{b_1} \varepsilon_{i1}^h$ and that by Slutsky symmetry $\frac{b_i}{b_1} \varepsilon_{i1}^h = \varepsilon_{1i}^h$, we have that:

$$\varepsilon_1(\alpha) = \frac{1}{1 - \alpha - \alpha \frac{b_2}{b_1}} \varepsilon_1^h + \frac{\alpha}{1 - \alpha - \alpha \frac{b_2}{b_1}} \sum_{i=3}^m \varepsilon_{1i}^h - \frac{1}{1 - \alpha - \alpha \frac{b_2}{b_1}} ((1-\alpha)\varepsilon_1^w b_1 - \alpha \varepsilon_2^w b_2).$$

Choosing the reference point so that $\frac{b_2}{b_1} \approx 0$, this simplifies to:

$$\varepsilon_1(\alpha) = \frac{1}{1-\alpha}\varepsilon_1^h + \frac{\alpha}{1-\alpha}\sum_{i=3}^{m}\varepsilon_{1i}^h - \varepsilon_1^w b_1,$$

where the first two terms provide the compensated elasticity of good 1 as a function of α and the third term provides the income effect.

Non-linear similarity

The inclusion of non-linear similarity is complicated by the additional goods, as this requires that the characteristics space is also m-dimensional, so that we need m measures to span the whole space that lies between B_1 and B_2. As an example, consider the case with 3 goods and 3 characteristics. Then good 2 can become more similar to good 1 by reducing the difference between the goods for any of the three characteristics. Hence, there is a cuboid for which each point is closer to B_1 than the original B_2. As this is a three-dimensional object, we cannot cover all its points with α and λ alone but need a third measure that is not a linear combination of the two. As it happens, it is convenient to add a second non-linear similarity measure (or multiple measures if $m > 3$) by allowing good 2 to move closer to good 3 (good $i \notin \{1, 2\}$) through β_3 (β_i) in a similar fashion as to good 1 through α. This is possible as C is invertible by assumption and thus B_1, B_2 and B_3 cannot lie on the same plane. Moreover, similarly to α that does not affect the price-elasticity of the third (i-th) good, β_3 (β_i) does not affect the price-elasticity of good 1. Hence, the own-price elasticity of good 1 is simply a function of α and λ also in this more general case, and we have:

$$\varepsilon_1(p, \alpha, \lambda) = \frac{1}{1-\alpha}\varepsilon_1^h(p, \lambda) + \frac{\alpha}{1-\alpha}\sum_{i=3}^{m}\varepsilon_{1i}^h(p, \lambda) - \varepsilon_1^w(p, \lambda)b_1.$$

Hence, as long as there is little substitution between good 1 and other goods than good 2 (i.e. $\sum_{i=3}^{m}\varepsilon_{1i}^h(p, \lambda) \approx 0$) and either the budget share of good 1 is small ($b_1 \approx 0$) or the income elasticity of good 1 is small (i.e. $\varepsilon_1^w(p, \lambda) \approx 0$), the own-price elasticity

$\varepsilon_1(p, \alpha, \lambda)$ is well approximated by:

$$\varepsilon_1(p, \alpha, \lambda) = \frac{1}{1-\alpha}\varepsilon_1^h(p, \lambda).$$

D APPENDIX - Welfare derivations

This appendix provides the formal derivations underlying the welfare analysis in Section 2.2. We first derive the marginal welfare effect of a tax change and the MVPF in the baseline two-good model with a separable externality from good 1 only. We then derive the optimal tax formula. Finally, we extend the analysis to n goods, to non-separable externalities, and to the case where both goods may generate externalities.

D.1 Derivation of the marginal welfare effect

We retain the setup of Section 2.2. There are two goods, $x = (x_1, x_2)'$, a continuum of consumers of mass one indexed by j, and the government levies a specific tax τ on good 1. Consumer j solves

$$\max_{x_j} \; u^j(Cx_j) \quad \text{s.t.} \quad (p_1 + \tau)x_{1j} + p_2 x_{2j} = m_j,$$

yielding demand functions $x_{1j}(\tau)$ and $x_{2j}(\tau)$, and indirect utility $v_j(\tau)$. The normalized indirect utility is $w_j(\tau) = v_j(\tau)/\lambda_j(\tau)$, where $\lambda_j(\tau)$ is the marginal utility of income. Consumption of good 1 generates a separable externality $e_j(c_{11}x_{1j})$ with $e_j' > 0$.

Social welfare is

$$W(\tau) = \int_j w_j(\tau)\, dj - \int_j e_j(c_{11}x_{1j}(\tau))\, dj + \mu\left[\tau x_1(\tau) + T - R\right], \tag{D.1}$$

where $x_1(\tau) = \int_j x_{1j}(\tau)\, dj$ is aggregate consumption of good 1, T is revenue from other sources, and μ is the shadow value of the government's revenue constraint.

Differentiating with respect to τ:

Consumer welfare term. By Roy's identity, $\frac{\partial v_j}{\partial \tau} = -\lambda_j x_{1j}$, so that $\frac{\partial w_j}{\partial \tau} = \frac{1}{\lambda_j}\frac{\partial v_j}{\partial \tau} = -x_{1j}$. Integrating over consumers:

$$\frac{d}{d\tau}\int_j w_j(\tau)\, dj = -\int_j x_{1j}\, dj = -x_1. \tag{D.2}$$

Externality term. By the chain rule:

$$\frac{d}{d\tau}\int_j e_j(c_{11}x_{1j}(\tau))\,dj = \int_j e_j'(c_{11}x_{1j})\cdot c_{11}\frac{\partial x_{1j}}{\partial \tau}\,dj. \tag{D.3}$$

Writing the individual demand response in elasticity form, $\frac{\partial x_{1j}}{\partial \tau} = \frac{\varepsilon_{1j}}{p_1+\tau}x_{1j}$, where $\varepsilon_{1j} = \frac{\partial x_{1j}}{\partial(p_1+\tau)}\frac{p_1+\tau}{x_{1j}}$ is the own-price elasticity of consumer j, we obtain

$$\frac{d}{d\tau}\int_j e_j(c_{11}x_{1j})\,dj = \frac{c_{11}}{p_1+\tau}\int_j e_j'\varepsilon_{1j}x_{1j}\,dj = \frac{e'c_{11}}{p_1+\tau}\varepsilon_1 x_1,$$

where we define the consumption-weighted marginal externality as

$$e' \equiv \frac{\int_j e_j'\varepsilon_{1j}x_{1j}\,dj}{\varepsilon_1 x_1}, \tag{D.4}$$

and the aggregate elasticity as $\varepsilon_1 = \int_j \varepsilon_{1j}\frac{x_{1j}}{x_1}\,dj$.

Revenue term. Differentiating $\tau x_1(\tau)$:

$$\frac{d}{d\tau}\left[\tau x_1(\tau)\right] = x_1 + \tau\frac{\partial x_1}{\partial \tau} = x_1 + \frac{\tau}{p_1+\tau}\varepsilon_1 x_1. \tag{D.5}$$

Combining. Substituting (D.2)–(D.5) into the derivative of (D.1):

$$\frac{dW}{d\tau} = -x_1 + \mu\left[x_1 + \frac{\tau}{p_1+\tau}\varepsilon_1 x_1\right] - \frac{e'c_{11}}{p_1+\tau}\varepsilon_1 x_1, \tag{D.6}$$

which is Equation (4) in the main text.

D.2 MVPF and optimal tax

MVPF. Rearranging Equation (D.6) as the ratio of the willingness to pay for the tax increase to the net cost to government, we define the MVPF as

$$\text{MVPF} = \frac{\text{WTP per €\ of mechanical revenue}}{\text{Net cost per €\ of mechanical revenue}} = \frac{x_1 + \frac{e'c_{11}}{p_1+\tau}\varepsilon_1 x_1}{x_1 + \frac{\tau}{p_1+\tau}\varepsilon_1 x_1}.$$

Dividing numerator and denominator by x_1 yields

$$\text{MVPF} = \frac{1 + \frac{e'c_{11}}{p_1+\tau}\varepsilon_1}{1 + \frac{\tau}{p_1+\tau}\varepsilon_1}, \tag{D.7}$$

which is Equation (6) in the main text.

The numerator includes one (the mechanical cost to consumers of paying one euro more in tax) plus the corrective benefit from reduced externalities. The denominator includes one (the mechanical revenue) plus the fiscal externality: the revenue change due to the behavioral response. When $\varepsilon_1 < 0$ (demand slopes downward), the fiscal externality is negative (raising the tax shrinks the base), so the denominator is less than one and the MVPF is pushed above one, all else equal. Meanwhile, when $e' > 0$, the corrective benefit adds to the numerator.

Optimal tax. At the optimum, $dW/d\tau = 0$, which from (D.6) requires

$$x_1 + \frac{e'c_{11}}{p_1 + \tau^*}\varepsilon_1 x_1 = \mu \left[x_1 + \frac{\tau^*}{p_1 + \tau^*}\varepsilon_1 x_1 \right].$$

Dividing both sides by x_1 and rearranging:

$$1 + \frac{e'c_{11}}{p_1 + \tau^*}\varepsilon_1 = \mu + \mu\frac{\tau^*}{p_1 + \tau^*}\varepsilon_1.$$

Isolating the terms with τ^*:

$$\mu\frac{\tau^*}{p_1 + \tau^*}\varepsilon_1 = 1 - \mu + \frac{e'c_{11}}{p_1 + \tau^*}\varepsilon_1.$$

Dividing by $\mu\varepsilon_1$ (with $\varepsilon_1 < 0$):

$$\frac{\tau^*}{p_1 + \tau^*} = \frac{1}{\mu}\frac{e'c_{11}}{p_1 + \tau^*} + \frac{\mu - 1}{\mu}\frac{1}{-\varepsilon_1}, \tag{D.8}$$

which is Equation (5) in the main text. The first term is the Pigouvian component and the second is the Ramsey (inverse elasticity) component.

Note that when $\mu = 1$ (first-best, lump-sum taxes available), the Ramsey term vanishes and we have $\tau^* = e'c_{11}$: the tax equals the marginal social damage, independent of the demand elasticity.

D.3 Extension to n goods

We now extend the welfare framework to a setting with n goods $x = (x_1, \ldots, x_n)'$, where the government levies specific taxes τ_k on goods $k = 1, \ldots, n$. Consumer j solves

$$\max_{x_j} \; u^j(Cx_j) \quad \text{s.t.} \quad \sum_{k=1}^{n} (p_k + \tau_k) x_{kj} = m_j.$$

Good k generates a separable externality $e_{kj}(c_{kk}x_{kj})$ with $e'_{kj} \geq 0$, where c_{kk} is the harm-relevant characteristic coefficient of good k. Some goods may have $e'_{kj} = 0$ for all j (no externality).

Social welfare is

$$W(\boldsymbol{\tau}) = \int_j w_j(\boldsymbol{\tau})\, dj - \sum_{k=1}^{n} \int_j e_{kj}(c_{kk}x_{kj}(\boldsymbol{\tau}))\, dj + \mu \left[\sum_{k=1}^{n} \tau_k x_k(\boldsymbol{\tau}) + T - R \right].$$

Differentiating with respect to τ_i for a particular good i, and following the same steps as in Section D.1:

$$\frac{\partial W}{\partial \tau_i} = -x_i + \mu \left[x_i + \sum_{k=1}^{n} \frac{\tau_k}{p_i + \tau_i} \varepsilon_{ki} x_k \right] - \sum_{k=1}^{n} \frac{e'_k c_{kk}}{p_i + \tau_i} \varepsilon_{ki} x_k, \tag{D.9}$$

where $\varepsilon_{ki} = \frac{\partial x_k}{\partial (p_i + \tau_i)} \frac{p_i + \tau_i}{x_k}$ is the elasticity of the aggregate demand for good k with respect to the consumer price of good i, and e'_k is the consumption-weighted marginal externality from good k (defined analogously to Equation D.4).

The cross-price elasticities ε_{ki} for $k \neq i$ capture the fact that changing the tax on good i also affects the consumption of all other goods, including other externality-generating goods. This has two consequences. First, the fiscal externality now includes the effect on the tax base of all goods, not just good i. Second, the corrective benefit (or cost) includes the change in all externalities, not just the one associated with good i.

Setting (D.9) to zero and solving yields the n-good analog of the optimal tax formula. In the special case of our two-good framework where only good 1 is taxed ($\tau_2 = 0$) and only good 1 generates an externality ($e_2' = 0$), cross-price effects on revenue and externalities vanish, and Equation (D.9) reduces to Equation (D.6).

The MVPF of increasing τ_i in the n-good setting is

$$\text{MVPF}_{\tau_i} = \frac{1 + \sum_{k=1}^{n} \frac{e_k' c_{kk}}{p_i + \tau_i} \frac{\varepsilon_{ki} x_k}{x_i}}{1 + \sum_{k=1}^{n} \frac{\tau_k}{p_i + \tau_i} \frac{\varepsilon_{ki} x_k}{x_i}}. \tag{D.10}$$

The additional cross-price terms in the numerator and denominator mean that the welfare evaluation of a tax change on good i must account for spillovers to other taxed goods and other externality-generating goods. In practice, when the cross-price elasticities between distant goods are small (which is the maintained assumption in our characteristics framework for goods that are not close substitutes), these additional terms are negligible and the two-good MVPF in Equation (D.7) provides a good approximation.

D.4 Non-separable externalities

In the main text, the externality is additively separable from individual utility in the social welfare function. Here we consider a case where the externality is non-separable: the social harm depends not only on consumption of good 1 but also on consumption of good 2. Specifically, let the externality be $e_j(c_{11}x_{1j}, c_{12}x_{2j})$, so that both the level of the harmful characteristic in good 1 and the level of some (possibly related) characteristic in good 2 jointly determine the social cost. For example, total sugar consumption may depend on both the sugary drink (good 1) and the alternative drink (good 2) if the alternative also contains some sugar.

Social welfare becomes

$$W(\tau) = \int_j w_j(\tau)\, dj - \int_j e_j(c_{11}x_{1j}(\tau),\, c_{12}x_{2j}(\tau))\, dj + \mu \left[\tau x_1(\tau) + T - R\right].$$

Differentiating with respect to τ:

$$\frac{dW}{d\tau} = -x_1 + \mu \left[x_1 + \frac{\tau}{p_1 + \tau} \varepsilon_1 x_1 \right] - \int_j \left[e'_{j,1} \cdot c_{11} \frac{\partial x_{1j}}{\partial \tau} + e'_{j,2} \cdot c_{12} \frac{\partial x_{2j}}{\partial \tau} \right] dj,$$

where $e'_{j,1} = \frac{\partial e_j}{\partial (c_{11} x_{1j})}$ and $e'_{j,2} = \frac{\partial e_j}{\partial (c_{12} x_{2j})}$ denote the partial derivatives of the externality function.

Writing the demand responses in elasticity form and defining aggregate cross-price elasticity $\varepsilon_{21} = \frac{\partial x_2}{\partial (p_1 + \tau)} \frac{p_1 + \tau}{x_2}$, this becomes

$$\frac{dW}{d\tau} = -x_1 + \mu \left[x_1 + \frac{\tau}{p_1 + \tau} \varepsilon_1 x_1 \right] - \frac{e'_1 c_{11}}{p_1 + \tau} \varepsilon_1 x_1 - \frac{e'_2 c_{12}}{p_1 + \tau} \varepsilon_{21} x_2, \tag{D.11}$$

where e'_1 and e'_2 are the appropriately weighted aggregate marginal externalities.

The additional term captures the externality effect of the tax through good 2. If taxing good 1 causes consumers to substitute toward good 2 ($\varepsilon_{21} > 0$), and good 2 also contributes to the externality ($e'_2 > 0$), then part of the corrective benefit of the tax is offset by increased harm through good 2. This "leakage" of the externality to the substitute good reduces the net corrective benefit and lowers the MVPF relative to the separable case.

The MVPF becomes

$$\text{MVPF} = \frac{1 + \frac{e'_1 c_{11}}{p_1 + \tau} \varepsilon_1 + \frac{e'_2 c_{12}}{p_1 + \tau} \frac{\varepsilon_{21} x_2}{x_1}}{1 + \frac{\tau}{p_1 + \tau} \varepsilon_1}. \tag{D.12}$$

The main intuition from the separable case carries over: the demand elasticity ε_1 remains central to the welfare evaluation, as it governs both the fiscal cost and the corrective benefit of the tax on good 1. The non-separability adds a cross-externality channel, but as long as the substitute good generates a smaller externality per unit consumed than the taxed good (i.e., $e'_2 c_{12} < e'_1 c_{11}$, as when sugar-free soda is less harmful than sugary soda), the tax retains a net corrective benefit.

D.5 General externalities: both goods taxed and harmful

Finally, we consider the most general two-good case where both goods may be taxed and both may generate externalities. The government levies taxes τ_1 on good 1 and τ_2 on good 2, and the externality depends on consumption of both goods through a separable structure: $e_j^1(c_{11}x_{1j}) + e_j^2(c_{22}x_{2j})$, where $e_j^{1\prime}, e_j^{2\prime} \geq 0$.

Consumer j solves

$$\max_{x_j} \; u^j(Cx_j) \quad \text{s.t.} \quad (p_1+\tau_1)x_{1j} + (p_2+\tau_2)x_{2j} = m_j.$$

Social welfare is

$$W(\tau_1,\tau_2) = \int_j w_j\, dj - \int_j \left[e_j^1(c_{11}x_{1j}) + e_j^2(c_{22}x_{2j})\right] dj + \mu\left[\tau_1 x_1 + \tau_2 x_2 + T - R\right].$$

Differentiating with respect to τ_1:

$$\frac{dW}{d\tau_1} = -x_1 + \mu\left[x_1 + \frac{\tau_1}{p_1+\tau_1}\varepsilon_{11}x_1 + \frac{\tau_2}{p_1+\tau_1}\varepsilon_{21}x_2\right] - \frac{e_1' c_{11}}{p_1+\tau_1}\varepsilon_{11}x_1 - \frac{e_2' c_{22}}{p_1+\tau_1}\varepsilon_{21}x_2, \tag{D.13}$$

where ε_{11} and ε_{21} are the own- and cross-price elasticities with respect to the consumer price of good 1, and e_1' and e_2' are the aggregate marginal externalities.

The MVPF of increasing τ_1 is

$$\text{MVPF}_{\tau_1} = \frac{1 + \frac{e_1' c_{11}}{p_1+\tau_1}\varepsilon_{11} + \frac{e_2' c_{22}}{p_1+\tau_1}\frac{\varepsilon_{21}x_2}{x_1}}{1 + \frac{\tau_1}{p_1+\tau_1}\varepsilon_{11} + \frac{\tau_2}{p_1+\tau_1}\frac{\varepsilon_{21}x_2}{x_1}}. \tag{D.14}$$

When good 2 is not taxed ($\tau_2 = 0$) and does not generate an externality ($e_2' = 0$), this reduces to the baseline MVPF in Equation (D.7). When good 2 is also taxed and harmful, the cross-price effects matter: substitution toward good 2 (when $\varepsilon_{21} > 0$) increases revenue from the tax on good 2 (raising the MVPF through the denominator) but also increases externalities from good 2 (lowering the MVPF through the numerator). The net effect depends on the relative magnitudes of τ_2 and $e_2' c_{22}$ for good 2.

This general formulation nests all the special cases considered in the main text and the preceding subsections. Regardless of which extension is considered, the demand elasticities, both own-price and cross-price, remain central to the welfare evaluation: they determine the magnitude of the behavioral responses that drive both the fiscal and corrective channels of the tax.