

Optimal sin taxation and market power

By MARTIN O'CONNELL AND KATE SMITH*

We study how market power impacts the efficiency and redistributive properties of sin taxation, with an empirical application to sugar-sweetened beverage taxation. We estimate an equilibrium model of the UK drinks market, which we embed in a tax design framework to solve for optimal sugar-sweetened beverage tax policy. Positive price-cost margins for drinks create inefficiencies that lower the optimal rate compared with a perfectly competitive setting. Since profits mainly accrue to the rich, this is partially mitigated under social preferences for equity. Overall, ignoring market power when setting tax policy leads to welfare gains 40% below those at the optimum.

JEL: D12, D43, D61, D62, H21, H23, L13

Keywords: externality, corrective tax, market power, profits, redistribution

Almost one-fifth of all consumer spending is undertaken in markets subject to taxes aimed, at least in part, at discouraging socially costly consumption.¹ Many of these markets are concentrated and likely to be characterized by firms that exercise market power. For instance, soft drink markets, recently the subject of new taxes in several jurisdictions (GFRP (2021)), are dominated by a small number of multi-product firms selling instantly recognizable and long-established brands viewed by consumers as imperfect substitutes.² The effective design of sin taxes requires balancing efficiency improvements and policymakers' redistributive goals. How best to do this depends on the competitive conditions in affected markets. Positive price-cost margins create allocative distortions, in addition to

* O'Connell: University of Wisconsin-Madison and the Institute for Fiscal Studies, Department of Economics, William H. Sewell Social Science Building. 1180 Observatory Dr, Madison, WI 53706, moconnell9@wisc.edu. Smith: London School of Economics and the Institute for Fiscal Studies, Houghton Street, London WC2A 2AE, k.e.smith2@lse.ac.uk. The authors would like to thank participants at seminars at Bristol, Brussels, Chicago Booth, INRAE, Michigan, Oxford, Warwick, UCL, UIUC, WUSTL and Wisconsin-Madison for many very helpful discussions. The authors gratefully acknowledge financial support from the Economic and Social Research Council (ESRC) under the Centre for the Microeconomic Analysis of Public Policy (CPP; grant number ES/T014334/1), under the Open Research Area (ORA; grant number ES/VO13513/1) and the British Academy under pf160093. Data supplied by Kantar FMCG Purchase Panel. The use of Kantar FMCG Purchase Panel data in this work does not imply the endorsement of Kantar FMCG Purchase Panel in relation to the interpretation or analysis of the data. All errors and omissions remain the responsibility of the authors.

¹Spending on alcohol, tobacco, soft drinks, fuel, and motoring (all of which are subject to some kind of excise duty in the UK – Levell, O'Connell and Smith (2016)) accounts for 19% of consumer spending recorded in the UK's consumer expenditure survey (Office for National Statistics, Department for Environment, Food and Rural Affairs (2007-2012)).

²More generally, a large industrial organization literature demonstrates that market power is a typical feature of many markets (see Einav and Levin (2010)), and there is recent evidence of its growing importance – for instance, see Syverson (2019) and De Loecker, Eeckhout and Unger (2020).

those associated with social costs of consumption; firms' strategic adjustment of margins affects the mapping between taxes, equilibrium prices and consumption; and concentration of profit ownership among the rich³ impacts the distributional consequences of reforms.

In this paper we show how market power affects the optimal design of sin taxes, and undertake a substantive empirical application to the taxation of sugar-sweetened beverages. We analyze a tax design framework, which highlights how market power impacts both the efficiency and redistributive effects of policy, and motivates our empirical approach. Harnessing detailed consumer level data, we estimate an equilibrium model of the UK market for drinks, and embed this into the tax problem to solve for optimal sugar-sweetened beverage tax policy. We show that allocative distortions associated with positive margins on drinks reduce the optimal sin tax rate, compared with a perfectly competitive setting, but this reduction is partially offset by the fact that profits accrue disproportionately to those with lower social welfare weights (foreign residents and the rich). Overall, ignoring market power leads to welfare gains that are 40% below those achieved by optimal policy.

Our tax design framework entails an environment with heterogeneous consumers and strategic firms. We study a government's choice over a linear commodity tax levied in a specific market of interest that comprises "sin" products, which generate social costs and alternatives, which do not create social costs. The products are supplied by multi-product firms that exercise market power, strategically reoptimize prices in response to tax policy, and earn positive profits, which are distributed to individuals in a potentially unequal manner. The government sets policy anticipating the strategic responses of firms,⁴ accounts for spillovers to other tax bases, and may place more weight on the welfare of the poor than the rich. We show how optimal policy balances: (i) reducing the social costs of consumption, (ii) exacerbating inefficiencies arising from the exercise of market power, (iii) fiscal externalities, and (iv) redistributive effects that arises both from inequality in consumption patterns, and inequality in profit holdings.

We apply our framework to study the taxation of sugar-sweetened beverages, a policy that is increasingly popular due, in part, to the large externalities that arise from the healthcare costs associated with treating obesity. The combination of asymmetrically differentiated products, multi-product firms and heterogeneous consumers means that the optimal tax expressions that we derive do not straightforwardly translate into externally valid elasticities. They depend instead on product-specific price-cost margins (which are typically unobservable, see Bresnahan (1989)), firm portfolio effects, and the full set of product level own- and cross-price elasticities. We therefore estimate a detailed model of consumer choice and firm pricing competition in the UK drinks market and embed this into

³See, for example, evidence in Cooper et al. (2016).

⁴This is also the case in Miravete, Seim and Thurk (2018), who show quantitatively how firms' pricing responses change the peak and shape of the Laffer curve in the liquor market.

our tax design framework. We model the discrete choice that consumers make between the many products in the market (or allocating all of their spending outside the market), and estimate the model using micro longitudinal data. We incorporate rich preference heterogeneity, including by income and a measure of total dietary sugar, which are important for capturing both redistributive and externality correcting aspects of tax policy. We exploit price changes that are agreed in advance by drinks firms and retailers, but that create differential variation across consumers, as a source of identifying price variation. We identify product level marginal costs by coupling our demand estimates with a game-theoretic pricing model.

Our demand and supply estimates show that drinks products are highly differentiated and, due to this and the multi-product nature of their portfolios, drinks firms exercise a substantial degree of market power. We find that the average Lerner index at observed prices is around 0.5 for both sugar-sweetened drinks and alternatives. Our estimates show that, in response to a change in the price of one sugar-sweetened drink, consumers are most willing to switch to a similar sugar-sweetened drink. However, in response to a rise in the price of all sugary drinks, substitution to alternative (non sugar-sweetened) drinks is substantial. Market power among alternative drinks is therefore relevant for sugar-sweetened beverage tax policy. Our model allows us to determine how firms adjust prices (and hence price-cost margins) in response to tax policy. We find that pass-through of a tax on sugar-sweetened drinks is slightly above 100%, and show that this is in line with (out-of-sample) observed price changes following the introduction of the UK's Soft Drinks Industry Levy, as well as with findings in other jurisdictions.

We use our estimated model to first solve for the optimal sin tax rate set by a planner concerned only with inefficiencies arising from externalities and market power, and that is indifferent to the distribution of incidence across individuals. Consumption of sugar-sweetened beverages is strongly linked to diet-related disease, which creates externalities through increased societal costs of funding both public and insurance-based health care (see Allcott, Lockwood and Taubinsky (2019b)). We use the best available evidence from the epidemiological literature to quantify these costs. The efficiency-maximizing tax rate balances the reduction in social costs of consumption – the more effective tax is at combating externality distortions, all else equal, the higher will be the optimal rate – against exacerbating distortions from the exercise of market power. If equilibrium price-cost margins for the sin products are high relative to untaxed alternative products, this will act to lower the optimal rate.⁵ We find that the efficiency-maximizing tax rate on sugar-sweetened beverages is 4.2 pence per 10 oz, which leads to a

⁵That the market power correction lowers the optimal rate if, on average, sin products have high margins *relative to alternatives* is consistent with the results in Kaplow (2021), who shows that in a Mirrleesian economy in which goods have fixed positive price-cost margins it is Pareto-improving to marginally lower/raise the tax rate on any single good that has an above/below average margin whilst off-setting the distributional effect through a distribution neutral adjustment to the labor tax schedule.

19% rise in their average price. If the government ignores all distortions from market power it would set a suboptimally high rate of 12 pence per 10 oz, which leads to efficiency *losses*.

We then consider optimal sin tax policy, assuming that the government does have redistributive concerns. In this case, both the share of profits collected by the government in corporate and dividend taxes, and the distribution of net-of-tax profits across individuals, affect optimal policy. We measure the allocation of corporate profits using information from national accounts and the distribution of dividend income across households. We find that under our baseline social preferences, the optimal sin tax rate on sugar-sweetened beverages is 6.0 pence per 10 oz,⁶ which is over 40% higher than the efficiency-maximizing rate. This increase reflects the net impact of two offsetting forces. On the one hand, sugar-sweetened beverage consumption is highest among relatively low income consumers, which acts to lower the optimal rate. On the other hand, post-tax profits are disproportionately owned by high income consumers (or flow overseas), which increases the progressivity of the tax, thus raising the optimal rate. Overall, the second effect dominates. We show that these results apply under a wide range of social preferences, including when the government places zero welfare weight on the post-tax profits flowing to individuals (in which case ignoring market power when setting policy would lead it to overshoot the optimal rate by 70% – with the resulting welfare gains 26% below optimal – due to spillovers to the corporate and dividend tax bases).

We also explore how the nature of social costs, including their overall size, degree of concentration among a specific group of individuals, and composition between externalities and internalities impact optimal policy. In addition, we simulate optimal policy under alternative tax instruments, including a multi-rate system and a sugar tax. In all cases market power has a first order impact on the efficiency and redistributive impacts of taxation, and ignoring when setting policy leads to substantial unrealized welfare gains.

Our paper contributes to a growing literature on quantifying the impacts of sugar-sweetened beverage taxation. This includes a set of papers that use the implementation of these taxes to estimate their effect on prices and quantities,⁷ and a set that use estimates of consumer demand based on periods and locations with no tax in place to simulate the introduction of taxes similar to those used in practice.⁸ Included in the second set of papers is Dubois, Griffith and O’Connell (2020), who use novel UK panel data to estimate individual-specific preference parameters for “on-the-go” soft drinks. We extend this work by esti-

⁶In practice, US and UK sugar-sweetened beverage taxes range from 7 to 15 pence per 10 oz.

⁷See, for instance, Bollinger and Sexton (2018) and Rojas and Wang (2017) who study the Berkeley tax, Seiler, Tuchman and Yao (2021) and Roberto et al. (2019) who study the Philadelphian tax, and Grogger (2017) who study the Mexican tax. For a full survey of the recent literature see Griffith et al. (2019).

⁸These papers include Bonnet and Réquillart (2013), Wang (2015), Harding and Lovenheim (2017), Chernozhukov, Hausman and Newey (2019).

inating consumer demand across the whole non-alcoholic drinks sector, including drinks brought into the home, modeling the supply-side of the market, and studying optimal policy design. Like us, Allcott, Lockwood and Taubinsky (2019a) ask: what is the optimal sugar-sweetened beverage tax? They develop a novel characterization of the optimal corrective commodity tax rate under general preference heterogeneity and an optimally set labor tax schedule, and incorporate consumer misoptimization into the corrective tax motive. Our work complements theirs by showing how market power impacts optimal sin tax policy.

Our work also adds to the emerging literature that uses empirically rich models of specific markets to evaluate how imperfect competition impacts public policy design. This literature includes work showing how the existence, and firms' strategic adjustment, of mark-ups impacts: revenue- and efficiency-maximizing taxation in alcohol markets (Conlon and Rao (2015), Miravete, Seim and Thurk (2018, 2020)); the effectiveness and optimal design of carbon abatement policy (Fowlie, Reguant and Ryan (2016), Preonas (2019)); and the design of demographic-targeted subsidies in health insurance markets (Tebaldi (2017), Polyakova and Ryan (2019), Einav, Finkelstein and Tebaldi (2019)). We contribute to this literature by quantifying how market power alters both the efficiency and redistributive effects, and hence optimal design, of sugar-sweetened beverage taxation. More generally, our results illustrate the importance of accounting for market power on commodity tax design.

The rest of the paper is structured as follows. In Section I we outline a tax design framework that highlights the key ingredients determining optimal policy. In Section II we summarize key characteristics of the UK drinks market. In Sections III and IV we present our empirical model of market demand and supply, and our estimates. In Section V we present our optimal tax results. In an Online Appendix we provide further details of the optimal tax problem, data, estimates and simulations.

I. Sin tax design

We begin by outlining a tax design framework, which serves to highlight the key determinants of optimal sin tax policy with market power, and which we use when solving for optimal sugar-sweetened beverage taxation. We consider a setting in which a government chooses tax policy in a market (for instance, the drinks market) with multiple products, including a set associated with social costs (e.g., sugar-sweetened beverages). We allow for the possibility that these products are sold by firms that exercise market power. The government sets the tax rates while accounting for interactions with other parts of the tax system, and balances distortions from social costs and market power with equity considerations.

A. Set-up

INDIVIDUALS. — There is a continuum of individuals (or consumers) indexed i . Individuals supply labor in a competitive labor market to generate pre-tax earnings, z^i , which are subject to a non-linear earnings tax, $\mathcal{T}(z^i)$. Each individual also potentially receives income arising from their holding of profits, which are generated by the sale of consumption goods. We denote total profits by Π and i 's share of profits by $\delta^i \geq 0$, where $\int_i \delta^i di = 1$. Individual profit holdings are subject to a (potentially non-linear) tax, $T_\Pi(\delta^i \Pi)$.⁹

Net income is spent on consumption goods: $\mathbf{x}_S^i = \{x_j^i\}_{j \in \mathcal{S}}$ is a quantity vector of “sin” products, belonging to the set \mathcal{S} , consumption of which potentially creates social costs; $\mathbf{x}_N^i = \{x_j^i\}_{j \in \mathcal{N}}$ is a quantity vector of products belonging to the set \mathcal{N} , which are in the same market as those in \mathcal{S} , but do not give rise to social costs; let $\mathcal{M} = \mathcal{S} \cup \mathcal{N}$ and $n(\mathcal{M}) = J$. x_O^i denotes the quantity of a composite consumption good that represents all goods outside of those in market \mathcal{M} . Consumers face the tax-inclusive price vector $\mathbf{p} = (\mathbf{p}_S, \mathbf{p}_N, 1)$, which embeds the normalization that the price of the composite good, x_O^i , is 1. The individual's budget constraint is $\sum_{j \in \mathcal{M}} p_j x_j^i + x_O^i = z^i - \mathcal{T}(z^i) + \delta^i \Pi - T_\Pi(\delta^i \Pi)$. We assume that the earnings tax is piece-wise linear, denote $\frac{dT}{dz^i} \equiv \tau_z^i$, and define virtual labor income by $G \equiv \tau_z^i z^i - \mathcal{T}(z^i)$. We denote the sum of virtual and unearned (profit) income by $Y^i \equiv G^i + \delta^i \Pi - T_\Pi(\delta^i \Pi)$.

Each individual chooses a bundle, $(\mathbf{x}_S^i, \mathbf{x}_N^i, x_O^i, z^i)$, to maximize utility, $U^i(\mathbf{x}_S^i, \mathbf{x}_N^i, x_O^i, z^i)$, subject to their budget constraint. Consumer i 's product demands are denoted by $x_j^i = x_j^i(\mathbf{p}, (1 - \tau_z^i), Y^i)$, earnings supply by $z^i = z^i(\mathbf{p}, (1 - \tau_z^i), Y^i)$, and indirect utility by $V^i = V^i(\mathbf{p}, (1 - \tau_z^i), Y^i)$. We denote the marginal utility of income by $\alpha^i \equiv \partial V^i / \partial Y^i$.

FIRMS. — Each product $j \in \mathcal{M}$ is produced by a single firm (firms can produce multiple products). We denote the market demand for product j by $X_j(\mathbf{p}_M) = \int_i x_j^i di$, and we denote the product's marginal cost by c_j . We consider a system of linear excise taxes that apply to the products in market \mathcal{M} , $\boldsymbol{\tau}_M$. Note that constant marginal costs means that if market \mathcal{M} were perfectly competitive, prices would shift mechanically with commodity taxes.

We allow for the possibility that firms exercise market power, meaning they can set price above marginal cost and face positive demand. In equilibrium (where all firms choose their strategies to maximize their profits), the tax-exclusive price for any product j ($\tilde{p}_j \equiv p_j - \tau_j$), can be written $\tilde{p}_j(\mathbf{c}_M, \boldsymbol{\Upsilon}_M; \boldsymbol{\tau}_M) = c_j + \mu_j(\mathbf{c}_M, \boldsymbol{\Upsilon}_M; \boldsymbol{\tau}_M)$, where μ_j denotes the equilibrium price-cost margin of product j .¹⁰

⁹ $T_\Pi(\delta^i \Pi)$ captures both corporate taxes and individual (e.g., dividend) taxation. For simplicity we write the earning and profit taxes as additively separable. This is unimportant for our analysis, which would not be materially affected by non-separabilities in $\mathcal{T}(\cdot)$ and $T_\Pi(\cdot)$.

¹⁰In this section we remain agnostic about the precise nature of the imperfect competition (e.g., whether products are offered by monopolist, competing oligopolists or colluding oligopolists). In Section III we assume the firms in the UK drinks market compete in a Nash-Bertrand pricing game.

For the composite good, which has aggregate demand $X_O = \int_i x_O^i di$, we assume that its price remains fixed in response to the introduction of an excise tax system in market \mathcal{M} , but we allow for the possibility of a non-zero price-cost margin, μ_O . Total variable profits in the economy are given by $\Pi \equiv \sum_{j \in \mathcal{M} \cup 0} \mu_j X_j$.¹¹

GOVERNMENT. — We consider a government that chooses a system of linear commodity taxes in market \mathcal{M} . The most general system entails a set of product specific taxes, $\tau_{\mathcal{M}} = \{\tau_j\}_{j \in \mathcal{M}}$. In practice, tax rates tend not to vary across disaggregate products, due to prohibitive implementation issues. We therefore focus on more constrained systems closer to those used in practice. We assume that when introducing the excise tax system the government does not change other parts of the tax system (in particular, the earnings and profit taxes).

Such tax reforms are often motivated by the existence of social costs associated with consumption. We allow for a budgetary externality associated with consumer i 's consumption of products in set $\mathcal{S} \subset \mathcal{M}$ that takes the form $\Phi^i \equiv \Phi^i(\mathbf{x}_{\mathcal{S}}^i)$, where $\Phi^i(\cdot)$ is weakly increasing in each of its arguments. We denote the marginal externality of consumer i 's consumption of good $j \in \mathcal{S}$ by $\frac{d\Phi^i}{dx_j^i} \equiv \phi_j^i$. Below we discuss the implications for tax policy of other forms of social costs, such as internalities.

The optimal choice of excise taxes requires the government to balance reducing inefficiencies associated with consumption externalities with the inefficiencies arising from the exercise of market power. It must also account for any spillovers to existing tax bases, and, depending on its preferences for equity, it may factor in distributional consequences of tax reform. We consider a government with the social welfare function:

$$(1) \quad W = \int_i \omega^i V^i + \lambda(T_D^i + T_{\Pi}(\delta^i \Pi) - \Phi^i) di,$$

where ω^i is the Pareto weight on consumer i , λ is the marginal value of government revenue, and where tax revenue raised from individual i is given by revenue from distortionary taxes:

$$T_D^i = \sum_{j \in \mathcal{M}} \tau_j x_j^i + \mathcal{T}(z^i)$$

and from the tax on their profit holdings, $T_{\Pi}(\delta^i \Pi)$.

Market power has important implications for tax design. First, the existence of positive margins distorts resource allocations. Second, the existence of positive

¹¹We hold fixed firms' entry and product design decisions. In the case of a sugar-sweetened beverage tax we believe this assumption is mild. In what follows we model demand and supply over the major brands that comprise 90% of the UK beverage market. These are long established and all available in both sugar-sweetened and diet varieties, meaning removing all sugar from a sugar-sweetened product to avoid the tax is equivalent to removing the product completely from the market – which our model suggests would lead to significant reductions in variable profits. In Section V.C we consider a tax levied directly on sugar content and allow for product redesign.

profits, depending on how they are distributed across individuals, may impact the distributional consequences of taxation. Third, as firms re-optimize their strategies in response to a tax change, the tax-exclusive prices of all products in the market may change in response to a change in the tax rate levied on any one product (with one implication of this being tax changes are not necessary shifted one-for-one to the products subject to the tax change).¹²

B. Optimal policy

We focus on the case where the government sets a single tax rate, τ_S , applied to the set of sin goods. This serves to clarify the key forces that determine optimal tax policy, facilitates comparison with existing sin tax results derived under perfect competition and is an interesting case as, in practice, governments often implement market specific excise tax systems that set a single rate.

The optimal sin tax rate, τ_S^* , satisfies the *implicit* function:

$$(2) \quad \tau_S^* = \underbrace{\bar{\phi} + \frac{\text{cov}(\phi_j^i, dx_j^i/d\tau_S)}{(1/n(\mathcal{S})) \times d\mathbb{X}_S/d\tau_S}}_{\text{externality correction}} - \underbrace{(\bar{\mu}_S - \bar{\mu}_N\Theta_N - \mu_O\Theta_O)}_{\text{market power correction}} + \underbrace{\frac{1}{d\mathbb{X}_S/d\tau_S} \left[\text{cov} \left(g^i, \sum_{j \in \mathcal{M}} x_j^i \rho_j - \delta^i (1 - \tau_{\Pi}^i) \frac{d\Pi}{d\tau_S} \right) \right]}_{\text{distributional concerns}} - \underbrace{\frac{d(\int_i \mathcal{T}(z^i) di)/d\tau_S}{d\mathbb{X}_S/d\tau_S}}_{\text{tax base erosion}}.$$

(see Appendix A). $\bar{\phi} \equiv \int_i \frac{1}{n(\mathcal{S})} \sum_{j \in \mathcal{S}} \phi_j^i di$ denotes the average marginal consumption externality in the population and $d\mathbb{X}_S/d\tau_S = \sum_{j \in \mathcal{S}} dX_j/d\tau_S$ is the impact of a marginal tax change on total consumption of the set of sin products. $\bar{\mu}_{\mathcal{X}} \equiv \sum_{j \in \mathcal{X}} \mu_k \frac{dX_j/d\tau_S}{\sum_{j' \in \mathcal{X}} dX_{j'}/d\tau_S}$ is the weighted average margin for products in set $\mathcal{X} = \{\mathcal{S}, \mathcal{N}, \mathcal{O}\}$, where the weights are each product's contribution to the marginal impact of the sin tax on equilibrium consumption of all products in that set. $\Theta_{\mathcal{X}} \equiv \frac{d\mathbb{X}_{\mathcal{X}}/d\tau_S}{d\mathbb{X}_{\mathcal{S}}/d\tau_S}$ is the fraction of reduced consumption of products in set \mathcal{S} diverted to those in $\mathcal{X} = \{\mathcal{N}, \mathcal{O}\}$ due to a marginal tax rise, and $\rho_j \equiv \frac{dp_j}{d\tau_S}$ is the impact of a marginal tax change on the equilibrium consumer price of product j . $g^i \equiv \frac{\omega^i \alpha^i}{\lambda}$ are social marginal welfare weights. Variation in g^i across individuals reflects the government's preferences for equity: a government with a preference for reducing inequality will assign low weights to the rich and high weights to the

¹²Related is previous work on optimal taxation and imperfect competition by Auerbach and Hines (2002), who study a setting with a representative consumer, homogeneous products and Cournot competition. Our set-up differs as we allow for heterogeneous consumers, differentiated products and social preferences for equity.

poor.

Equation (2) expresses the optimal tax rate as a function of wedges from non-government (externality and market power) distortions, tax derivatives (for quantities, prices and profits) and government distributional preferences. It comprises four components, three reflecting efficiency and one equity considerations. Although these components are written additively, it should be noted that the terms on the right hand side of the expression (including product-level margins) depend on the tax rate, and that market power influences each component of the formula, because the strategic responses of firms are key in driving the quantity, price and profit tax derivatives. We describe each component of the formula in turn.

EXTERNALITY CORRECTING COMPONENT. — This equals the average marginal externality across consumers and sin products, plus an adjustment that reflects the covariance between the consumer-product specific marginal externality and the sensitivity of the individual's consumption of the product to a change in the sin tax rate. This covariance captures how effective the tax is at reducing the most socially costly consumption. The more a tax rise reduces consumption by consumers and/or of products associated with high marginal externalities, the better targeted it will be at the most externality generating consumption and the higher will be the optimal tax rate.¹³ This mirrors the logic in the optimal externality correcting tax rate with heterogeneous externalities derived in Diamond (1973). However, an important difference is that the sensitivity of consumption to changes in the sin tax rate also depends on the equilibrium pricing response of firms; i.e.,

$$\frac{dx_j^i}{d\tau_S} = \sum_{j' \in \mathcal{M}} \frac{\partial x_j^i}{\partial p_{j'}} \frac{dp_{j'}}{d\tau_S}.$$

MARKET POWER CORRECTING COMPONENT. — All else equal, higher equilibrium margins on the products in set \mathcal{S} act to lower the optimal tax rate. This reflects the argument made by Buchanan (1969), who pointed out the appropriate tax rate on an externality-generating monopolist lies below the full Pigouvian rate. However, if taxing products in set \mathcal{S} induces an increase in consumption of other products also supplied non-competitively, then distortions arising from the exercise of market power on these alternatives, all else equal, act to raise the optimal tax rate. The strength of this effect depends on the weighted average equilibrium margin on these non-taxed alternatives and the extent to which a marginal tax rise diverts equilibrium consumption towards them. For products in the set \mathcal{N} , these two forces are captured by the terms $\bar{\mu}_{\mathcal{N}}$ and $\Theta_{\mathcal{N}}$ (analogous expressions capture the influence of any market power distortions from the numeraire good).

¹³Note, $dX_{\mathcal{S}}/d\tau_S$ will generally be negative; an increase in the sin tax rate will lower equilibrium consumption of those products. If $\text{cov}(\phi_j^i, dx_j^i/d\tau_S)$ is negative, so a tax rise tends to achieve relatively large consumption reductions among products/consumers with large marginal externalities, the externality correction will exceed $\bar{\phi}$. Hence, all else equal, the more the total reduction in sin good consumption is concentrated among the most socially costly consumption, the higher will the externality correction component of the optimal tax rate.

DISTRIBUTIONAL CONCERNS. — These enter equation (2) through the covariance of consumers' social marginal welfare weights with the reduction in their utility resulting from a marginal increase in the tax rate. The more a marginal tax rise results in relatively large utility falls among those with high social marginal welfare weights (low income consumers when the planner has a preference for equity), the lower will be the optimal tax rate. This covariance term is scaled by the marginal effect of the tax on total consumption of the sin products. Hence, the strength of the distributional channel depends on how responsive equilibrium consumption of the set of sin products is to a marginal sin tax rise: the more sensitive is equilibrium consumption, the smaller the impact of equity considerations on the optimal tax rate.

To highlight how market power distortions interact with distributional concerns, it is informative to consider this term under perfect competition: $\frac{1}{dX_S/d\tau_S}$ $\text{Cov}\left(g^i, \sum_{j \in S} x_j^i\right)$. The reduction in a consumer's utility due to a marginal tax rise equals their total consumption on the taxed sin goods. All else equal, the more that those with high social marginal welfare weights consume a large quantity of the set of sin goods, the lower is the optimal tax rate on them.¹⁴ Imperfect competition has two consequences for the utility impact of a marginal tax rate increase, and hence on the distributional component of equation (2). First, consumption of the sin products is replaced by consumption of all products in market \mathcal{M} , weighted by the marginal impact of tax on each product's equilibrium price. If firms hold their tax-exclusive prices fixed in response to tax changes, then this term collapses back to total consumption of the sin goods. Second, the impact of the tax on the size of individuals' net-of-tax profit holdings also matters. If the tax leads to a reduction in profits, and profit holdings are disproportionately held by those with low social marginal welfare weights (the wealthy), this will act to make the tax more progressive and will increase the optimal rate.

TAX BASE EROSION. — This term arises because we assume that the government holds fixed the earnings tax schedule. All else equal, the more that a marginal increase in the sin tax leads to a reduction in labor tax revenue, the lower will be the optimal rate. Whether this term (and hence the loss from not re-optimizing the earnings tax alongside introducing the sin tax) is large or small is context dependent. To highlight what drives this term we assume that income effects on labor supply are negligible (see Saez, Slemrod and Giertz (2012) for empirical support of this), which allows us to re-write $d(\int_i \mathcal{T}(z^i) di)/d\tau_S$ as

$$\frac{d(\int_i \mathcal{T}(z^i) di)}{d\tau_S} = \int_i \frac{\tau_z^i}{1 - \tau_z^i} \zeta_z^i \sum_{j \in \mathcal{M}} \xi_j^i x_j^i \rho_j di,$$

¹⁴Allcott, Lockwood and Taubinsky (2019a) show that if the government also optimizes the earnings tax schedule, what matters is the cross-sectional correlation in social marginal welfare weights and consumption of the taxed goods *net of income effects*.

where ζ_z^i is the individual (compensated) elasticity of taxable earnings and ξ_j^i is the individual elasticity of demand for product j with respect to earnings (see Appendix A). This expression highlights that a key determinant of the tax base erosion component of equation (2) is the strength of income effects for products in market \mathcal{M} . In Appendix B we provide empirical evidence that these are very small in the context of the UK drinks market, nonetheless, they may be significant in other settings.

OTHER FORMS OF SOCIAL COSTS. — Sin taxes are often partially motivated by internalities – costs that individuals impose on themselves that they ignore when making choices. Internalities affect social welfare through individual utilities, in contrast to a budgetary externality, which affects welfare through its impact on the government’s budget. Suppose individual i ’s consumption of sin products gives rise to an externality, $I^i = I(\mathbf{x}_S^i)$, which impacts their normative but not decision utility, and let $\psi_j^i \equiv \frac{\partial V^i / \partial I^i}{\alpha^i} \frac{\partial I^i}{\partial x_j^i}$ denote the monetary cost the individual imposes on themselves per additional unit of consumption of sin good j due to this bias. In this case, there are two modifications to the optimal policy condition (2) – see Appendix A. First, the externality correcting term will now reflect externality and externality correction; ϕ_j^i is replaced with $\phi_j^i + \psi_j^i$. Second, the distributional concerns component contains an additional term capturing the covariance of externality changes and social marginal welfare weights; all else equal, for a government with preferences for equity, the more externality reductions resulting from a tax rise are concentrated among the relatively poor, the more valuable is a given fall in the average externality and hence the higher is the optimal tax rate. The impact of these considerations are studied in detail in Allcott, Lockwood and Taubinsky (2019a).

C. Discussion of empirical implementation

A common approach to empirical tax analysis is to write the expression of interest in terms of a set of externally valid elasticities, or sufficient statistics. An advantage of this approach is that the elasticities can be estimated using quasi-experimental variation, with transparent identification arguments (Chetty (2009)). However, an important restriction of the application of sufficient statistics to optimal tax formulae (as opposed to marginal tax reforms) is that they require implicitly assuming an iso-elastic preference model – see Kleven (2021).

There are two challenges with implementing a sufficient statistics approach in our context. First, the tax derivatives do not straightforwardly map into price elasticities. For a number of the derivatives this could potentially be overcome by making simplifying assumptions about tax pass-through.¹⁵ Alternatively, one could use data covering the introduction of a new (or change in an existing)

¹⁵For instance, in general, the impact of a marginal tax rise on consumption of the set of sin products

tax to directly estimate the tax derivatives (policy elasticities in the language of Hendren (2016)). The second challenge is that the optimal tax formula depends on equilibrium product-level price-cost margins. Marginal costs, and hence margins, are typically not straightforwardly observable in economic data. However, they can be inferred based on a profit-maximizing model of the firms operating in the market, coupled with estimates of the own- and cross-price elasticities of all the products in the market.

Our approach is therefore to specify and estimate an equilibrium model of the market of interest (the UK drinks market). This enables us to simulate the impact of an arbitrary tax policy on equilibrium margins, consumption and profits. To validate the model, we compare, where possible, its predictions to existing estimates of relevant elasticities, as well as to evidence from the introduction of the UK’s sugary drinks tax. We embed the equilibrium model into the tax problem and consider optimal policy under different government preferences for redistribution.

II. The drinks market

Sugar-sweetened beverage taxation is a natural setting in which to study how taxes designed to reduce social costs interact with market power. In many jurisdictions, taxes on drinks are explicitly motivated as a tool for improving public health, in part due to substantial social costs associated with their consumption. We discuss the nature and measurement of these costs in Section V. It is also the case that the market is concentrated and comprises a set of highly recognizable branded products. It is likely therefore that firms exercise considerable market power.

A. Data

We model behavior in the UK market for non-alcoholic drinks. This market includes carbonated drinks (often referred to as sodas), fruit concentrates, and sports and energy drinks. We refer collectively to these as soft drinks. Soft drinks brands usually come in sugar-sweetened and artificially sweetened (i.e., diet) varieties. Sugar-sweetened beverage taxes typically apply to sugar-sweetened varieties of soft drinks. The market also includes pure fruit juices and flavored milk. We use micro data on the drinks purchases of a sample of consumers living in Great Britain collected by the market research firm Kantar. The data contain information on household level purchases for “at-home” consumption from the Kantar Worldpanel, as well as purchases made by individuals for “on-the-go” consumption from Kantar On-The-Go Survey (Kantar (2008-2012)). Together “at-home”

takes the form $\frac{d\mathbb{X}_S}{d\tau_S} = \sum_{j \in S} \sum_{j' \in \mathcal{M}} \frac{\partial X_j}{\partial p_{j'}} \frac{dp_{j'}}{d\tau_S}$. However, under the assumption of fixed tax pass-through (denoted ρ) across products in S and fixed tax-exclusive prices for other goods, this collapses to $\rho \times \nabla \mathbb{X}_S$ where we use $\nabla \mathbb{X}_S$ to denote the marginal impact of consumption goods in S with respect to a marginal price rise for all these products.

and “on-the-go” consumption account for over 90% of drinks consumption by volume.¹⁶

There are 30,405 households in the at-home sample who record, by means of a barcode scanner, all grocery purchases made and brought into the home. The data are broadly representative of British households (in Appendix B we compare them with the nationally representative Living Cost and Food Survey) and cover 2008 to 2012. Individuals in the on-the-go data record all purchases they make from shops and vending machines for out-of-home consumption using a cell phone app. The data cover 2010 to 2012 and comprise 2,862 individuals (aged 13 and upwards) randomly drawn from the Worldpanel households. In both datasets, we observe households/individuals over many months. The data contain detailed information – including brand, flavor, size and nutrient composition – on the UPCs (barcodes) purchased, the store in which the purchase took place, and transaction level prices.

B. Consumers and purchasing patterns

We use the term consumer to refer to households in the at-home segment, and individuals in the on-the-go segment. Figure 1 highlights variation in soft drinks purchases across two key dimensions. Panel (a) shows that consumers that get a high fraction of their *total* dietary calories from added sugar purchase significantly more sugar-sweetened beverages than other consumers. Policymakers have typically focused on changing the behavior of consumers with dietary sugar above a particular threshold, due to elevated health risks (e.g., the World Health Organization (2015) advice focuses primarily on those with dietary sugar above 10%). The more that a sugar-sweetened beverage tax is able to achieve large consumption reductions among those consumers that create relatively high marginal externalities through their sugar-sweetened beverage intake, the more effective it will be at reducing externality distortions.

Panel (b) shows that there is a negative cross-sectional correlation between sugar-sweetened beverage consumption and equivalized household income – richer households consume less sugar-sweetened beverages, and therefore have consumption baskets less exposed to a sugar-sweetened beverage tax, than lower income consumers. The extent to which this is driven by preference heterogeneity (correlated with income) versus causal income effects determines how much sugar-sweetened beverage taxation induces labor supply distortions (captured by the tax base erosion term in equation (2)), and hence impacts optimal policy. In Appendix B we show that after removing consumer fixed effects (and hence relying on within household income transitions to estimate any income effects), the consumption gradient in equivalized income flattens completely. In our demand model we control flexibly for equivalized income, and in our optimal tax analysis,

¹⁶The remainder occurs in restaurants and bars, which are not covered by our data. Numbers are computed using the Living Cost and Food Survey.

we treat variation in drinks demand across the income distribution as preference heterogeneity.

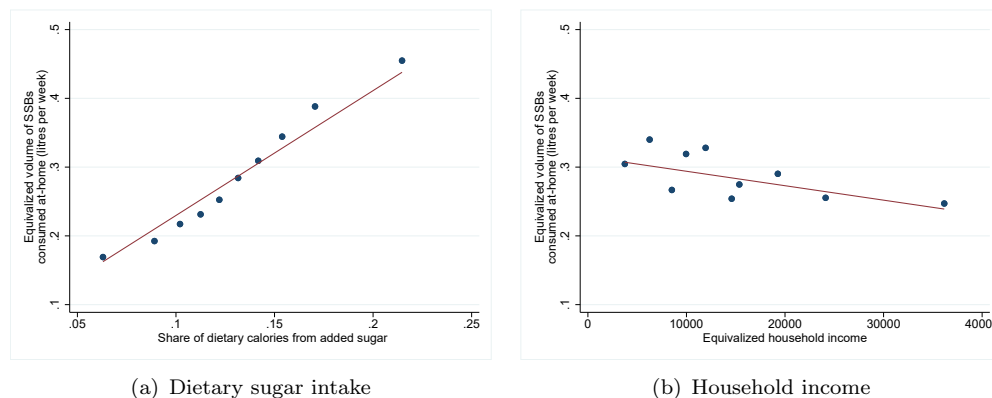


FIGURE 1. VARIATION IN VOLUME OF SUGAR-SWEETENED BEVERAGES CONSUMED AT-HOME

Note: The left hand panel shows mean volume of sugar-sweetened beverages purchased per person per week and consumed at home by deciles of the share of dietary calories from added sugar (from food consumed at home). The right hand panel shows mean volume of sugar-sweetened beverages purchased per person per week and consumed at home by deciles of equivalized (using the OECD-equivalence scale) household income. Analogous figures for the sugar from soft drinks consumed on-the-go are shown in Appendix B.

C. Firm and brands

The drinks market is highly concentrated. In Figure 2 we show the cumulative market share of the ten largest UK producers. The two largest firms (Coca Cola Enterprises and Britvic have a combined market share of almost 60%).¹⁷

Firms each typically own several separate brands. For instance, Coca Cola Enterprises most popular brand is Coca Cola, but it also owns 9 other brands with market share of at least 1%. Each soft drinks brand is typically available in sugar-sweetened (“regular”) and artificially sweetened (“diet” and/or “zero”) variants. Each brand-variant is available in multiple pack-sizes. In our equilibrium model of the market we focus on the set of brands with more than a 1% market share in either the at-home or on-the-go segment, as well as the main fruit juice and flavored milk brands, which together comprise over 75% of total spending on non-alcoholic drinks. Table B.3 in Appendix B lists brand-variants and the number of sizes they are available in.

¹⁷Drinks producers are known in the industry as bottlers. They buy concentrate from upstream firms (e.g., Coca Cola Enterprises obtains concentrate from The Coca Cola Company) and use this as an input to produce soft drinks products (see Luco and Marshall (2020)). For modeling firm behaviour in the UK drinks market, it is the bottlers – all of which are national – who are the relevant agents: as well as producing the products, they are responsible for negotiating product shelf prices and placement with retailers, and for promotional activity (see Competition Commission (2013)).

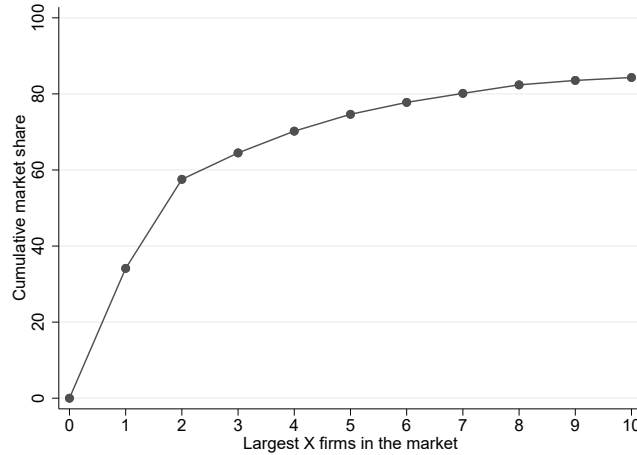


FIGURE 2. MARKET SHARE OF THE LARGEST FIRMS IN THE DRINKS MARKET

Note: The line shows the cumulative market share for the X largest firms in the market, where X is shown on the horizontal axes. Market shares are shown for 2012.

D. Drinks firm-retailer relations

Drinks firms do not sell directly to consumers, rather retailers act as intermediaries between drinks firms and consumers. The majority of expenditure is undertaken in national grocery chains (see Table B.4 in Appendix B). In the UK the main grocery chains set prices nationally (see Competition Commission (2000)), which means that if we observe a product's price in one store, we know the price faced by consumers in other stores belonging to the same retailer chain. Retailers typically offer all brand-varieties, though the available pack sizes can vary.

We do not directly observe the contracting relationship between the drinks firms and retailers. However, a 2013 report into the soft drinks market by the UK competition authority provides evidence on the nature of these relations (see Competition Commission (2013)). They cite evidence that annual bilateral "Joint Business Plans" are agreed between a drinks firm and retailer setting out wholesale prices, payments related to product visibility, recommended retail prices, and agreements on the number, type and timings of promotions. This evidence of non-linear contracting suggests drinks firms and retailers avoid double marginalization. We therefore treat drinks firms as (effectively) setting final consumer prices, an outcome consistent with optimal non-linear contracting – see Villas-Boas (2007) and Bonnet and Dubois (2010). We also exploit the fact that the promotions are agreed on in advance (and are not coordinated across retailers) as a useful identifying source of price variation (see Section III.C).

III. Equilibrium model of the drinks market

We estimate a model of consumer demand in the drinks market using a discrete choice framework in which consumer preferences are defined over product characteristics (Gorman (1980), Lancaster (1971), Berry, Levinsohn and Pakes (1995)). This enables us to model demand for the many differentiated products in the market, while incorporating rich preference heterogeneity, including by total dietary sugar and income. We identify product level marginal costs by combining the demand estimates with the equilibrium conditions from an oligopoly pricing game (Berry (1994), Nevo (2001)). The estimates of the primitives of demand and supply enable us to simulate the impact of tax policy on equilibrium quantities and prices, and hence consumer utilities and profits. Appendix C provides additional details on the demand specification.

A. Consumer demand

CHOICE PROBLEM. — We model which, if any, drink product a consumer (indexed i) chooses on a “choice occasion”, where choice occasion refers to a week in which a household purchases groceries in the at-home segment, or a day on which an individual buys a cold beverage (including bottled water) in the on-the-go segment. We treat the decisions that households make in the at-home segment and individuals make in the on-the-go segment separately, allowing for all preferences to vary freely with each type of choice situation. In Appendix C we provide evidence that recent purchases of drinks by a household in the at-home segment do not influence either the propensity to buy or quantity purchased by household members in the on-the-go segment. Choice in the on-the-go segment is between single portion size of products (e.g., 330ml cans and 500ml bottles); choice in the at-home segment is between multi-portion sizes. For notational parsimony we suppress a market segment index.

We index the drinks products by $j = \{1, \dots, J\}$. Products vary by brand, indexed by $b = \{1, \dots, B\}$, whether or not they contain sugar (for instance, the brand Coke is available in Regular, Diet and Zero variants), and their size, indexed by $s = \{1, \dots, S\}$. Brand-variants can be purchased in different sizes for two reasons: (i) the availability of different pack sizes (or UPCs), and (ii) the purchase of multiple packs. For each brand-variant we define sizes as the set of available pack sizes and the most common multiple pack purchases of UPCs. The consumer chooses between the available drinks products and choosing not to buy a drink, which we denote by $j = 0$. On around 42% of at-home and 60% of on-the-go choice occasions, a household purchases a drink (i.e., $j > 0$).¹⁸

¹⁸Consumers are sometimes observed purchasing multiple (typically) two brand-variants on a single choice occasion. On 40% (10%) of occasions in which a consumer chooses a drink in the at-home (on-the-go) segment, multiple are chosen. In this case, we randomly sample one, assuming that, conditional on consumer specific preferences, these purchases are independent, e.g., because they are bought for different household members.

We take as given the consumer's decision over which retailer (indexed r) to shop with. Choice sets are retailer specific (and denoted Ω_r) due to some variation in available pack-sizes by retailer.

Consumer i in period t , with total period budget y_{it} , solves the utility maximization problem:

$$(3) \quad V(y_{it}, \mathbf{p}_{rt}, \mathbf{x}_t, \epsilon_{it}; \mathbf{v}_i) = \max_{j \in \{\Omega_r \cup 0\}} \nu(y_{it} - p_{jrt}, \mathbf{x}_{jt}; \mathbf{v}_i) + \epsilon_{ijt}.$$

where $\mathbf{p}_{rt} = (\mathbf{p}_{1rt}, \dots, \mathbf{p}_{Jrt})$ is the price vector faced by the consumer, \mathbf{x}_{jt} are other characteristics of product j , and $\mathbf{x}_t = (\mathbf{x}_{1t}, \dots, \mathbf{x}_{Jt})$ (note $p_0 = 0$ and $\mathbf{x}_{0t} = 0$); \mathbf{v}_i is a vector of consumer level preference parameters; and $\epsilon_{it} = (\epsilon_{i0t}, \epsilon_{i1t}, \dots, \epsilon_{iJt})$ is a vector of idiosyncratic shocks.

The function $\nu(\cdot)$ captures the payoff the consumer gets from selecting option j . Its first argument, $y_{it} - p_{jrt}$, is spending on the numeraire good, i.e., spending outside the drinks market. We assume that preferences are quasi-linear, so $y_{it} - p_{jrt}$ enters $\nu(\cdot)$ linearly and y_{it} differences out when the consumer compares different options; we therefore suppress the dependency of $\nu(\cdot)$ on y_{it} . An implication of quasi-linearity is that a change in the price of any drinks product does not induce an income effect. Given the small share of total consumer expenditure allocated to drinks products, this is a mild assumption. However, we do allow equalized household income to shift consumer preferences, v_i . This enables our model to capture how demands vary across consumers with different incomes.

We assume that ϵ_{ijt} is distributed i.i.d. type I extreme value. Under this assumption the probability that consumer i selects product j in period t , conditional on prices, product characteristics and preferences, is given by:

$$(4) \quad \sigma_j(\mathbf{p}_{rt}, \mathbf{x}_t; \mathbf{v}_i) = \frac{\exp(\nu(p_{jrt}, \mathbf{x}_{jt}; \mathbf{v}_i))}{1 + \sum_{j' \in \Omega_r} \exp(\nu(p_{j'rt}, \mathbf{x}_{j't}; \mathbf{v}_i))},$$

and the consumer's expected utility is given by:

$$(5) \quad v(\mathbf{p}_{rt}, \mathbf{x}_t; \mathbf{v}_i) = \ln \sum_{j \in \Omega_r} \exp\{\nu(p_{jrt}, \mathbf{x}_{jt}; \mathbf{v}_i)\} + C,$$

where C is a constant of integration.

UTILITY SPECIFICATION. — We allow for rich preference heterogeneity, with both observed and unobserved consumer characteristics. This is important for two reasons. First, it is well established that the inclusion of rich, and, in particular, unobserved preference heterogeneity is crucial in enabling models of this type to recover realistic patterns of consumer substitution across products (see Berry, Levinsohn and Pakes (1995), Nevo (2001)). Second, in our application it is important that we capture variation in preferences across different consumers

that is relevant for tax policy. This includes variation across consumers whose consumption is likely to create different externalities at the margin, and across consumers to which a government may assign different social marginal welfare weights.

We therefore partition consumers into groups (indexed d) based on two variables. First, whether their total share of dietary calories (measured in the preceding year) is below 10%, between 10 and 15% and in excess of 15%. Second, we allow for at-home preferences to vary by whether the household contains children and on-the-go preferences to vary by whether the individual is younger or older than 30. We allow all preference parameters to vary by these groups; see Table C.1 in Appendix C for details. We also allow preferences over key product attributes, including price, to vary with equivalized household income (which we denote \tilde{y}_i).

The payoff function $\nu(\cdot)$ for consumer i belonging to consumer group $d(i)$ and for product j belonging to brand $b(j)$ and of size $s(j)$ takes the form:

$$\nu(\cdot) = -\alpha_{i0}p_{jrt} + \sum_{k>0}^K \alpha_{ik}x_{jk} + \zeta_{d(i)b(j)s(j)rt},$$

where, for product attribute, $k = 0, \dots, K$:

$$\alpha_{ik} = \bar{\alpha}_{d(i)k} + \alpha_{d(i)k}^{\tilde{y}} \tilde{y}_i + \sigma_{d(i)k}^{\alpha} \eta_{ik},$$

$\zeta_{d(i)b(j)s(j)rt}$ denotes an unobserved brand-size attribute (which may vary by consumer group, retailer and time) and η_{ik} is a standard normal random variable.¹⁹

$\bar{\alpha}_{d(i)k}$ denotes the baseline preference for product attribute k among consumer group d , $\alpha_{d(i)k}^{\tilde{y}}$ captures how preferences for the attribute vary across consumers with different equivalized household incomes, and $\sigma_{d(i)k}^{\alpha}$ captures the dispersion in unobserved preferences for the attribute. In addition to price, product attributes include sugar content, drink types (cola, lemonade, pure fruit juice etc.) size and advertising.²⁰ We allow preferences over price, sugar, branded soft drinks, non-branded drinks and pure fruit juice to vary with equivalized income, and for unobserved heterogeneity in preferences for drinks (relative to the outside option), price, sugar, size, cola, lemonade, non-branded drinks and fruit juice. We allow for correlation between preferences for sugar and for soft drinks (relative to the outside option). We allow both the baseline and dispersion parameters to vary across consumer groups d , so the overall random coefficient distribution is a

¹⁹For price we assume the coefficient is log-normally distributed.

²⁰We measure monthly advertising expenditure across TV, radio and online in the AC Nielsen Advertising Digest (AC Nielsen (2007-2012)). We compute product specific stocks based on a monthly depreciation rate of 0.8. This is similar to the rate used in Dubois, Griffith and O'Connell (2018) on similar data in the potato chips market.

flexible mixture of normal distributions.

We decompose the unobserved product attribute, $\zeta_{d(i)b(j)s(j)rt}$, into a set of detailed fixed effects:

$$\zeta_{d(i)b(j)s(j)rt} = \xi_{d(i)b(j)s(j)}^{(1)} + \xi_{d(i)b(j)t}^{(2)} + \xi_{d(i)s(j)t}^{(3)} + \xi_{d(i)s(j)r}^{(4)} + \xi_{d(i)b(j)r}^{(5)}$$

All the fixed effects are consumer group specific. They include brand-size, brand-time (year-quarter), size-time (year-quarter), brand-retailer and size-retailer effects. These control for shocks to demand that may be correlated with price setting.

B. Supply model

We model drinks firms as setting prices in a simultaneous move Nash-Bertrand game.²¹ We do not explicitly model drinks firms-retailer relationships, but, based on the evidence of non-linear contracting in these relations, we assume manufacturers set consumer prices, which is consistent with efficient contracting. Let $\mathbf{p}_m = (p_{1m}, \dots, p_{Jm})$ denote the prices that drinks firms set in market (year) m .²²

Market demand for product j is given by:

$$q_{jm}(\mathbf{p}_m) = \int_i \sigma_j(\mathbf{p}_m, \mathbf{x}_m; \mathbf{v}_i) dF(\mathbf{v}) M_m,$$

where M_m denotes the potential size of the market. We denote the marginal cost of product j in market m as c_{jm} .

We index the drinks firms by $f = (1, \dots, F)$ and denote the set of products owned by firm f by \mathcal{J}_f . Firm f 's total variable profits in market m are

$$(6) \quad \Pi_{fm}(\mathbf{p}_m) = \sum_{j \in \mathcal{J}_f} (p_{jm} - c_{jm}) q_{jm}(\mathbf{p}_m).$$

Under Nash-Bertrand competition, the equilibrium prices satisfy the set of first order conditions: $\forall f$ and $\forall j \in \mathcal{J}_f$,

$$(7) \quad q_{jm}(\mathbf{p}_m) + \sum_{j' \in \mathcal{J}_f} (p_{j'm} - c_{j'm}) \frac{\partial q_{j'm}(\mathbf{p}_m)}{\partial p_{jm}} = 0.$$

²¹Around one-fifth of the market consists of “store brands”. These are no-frills, low priced alternatives to the branded products. We treat these as a competitive fringe, with (tax-exclusive) prices that remain fixed.

²²In practice, for a given product-year a drinks firm and retailer agree on a base price \bar{p} and a sale price p_S , with the former applying ρ proportion of weeks (see Section II.D). Instead of modeling choice over (\bar{p}, p_S, ρ) , we model choice over $p = (1 - \rho)\bar{p} + \rho p_S$. This average price exhibits little variation across retailers. Cross-retailer variation in the price of a given product at a point in time is driven by non-synchronization of sales. Hence, we specify the relationship between prices in the supply game, p_{jm} , and those faced by consumers in retailer r week $t \in m$ as $p_{jrt} = p_{jm} + e_{jrt}$, where $\mathbb{E}[e_{jrt} | (j, m)] = 0$.

From this system of equations we can solve for the implied marginal cost, c_{jm} , and hence the equilibrium price-cost margin, $\mu_{jm} = p_{jm} - c_{jm}$, for each product in each market. For any set of product taxes we can use the system of equations (7), replacing c_{jm} with $c_{jm} + \tau_j$, to solve for counterfactual equilibrium prices p'_{jm} and margins $\mu'_{jm} = p'_{jm} - \tau_j - c_{jm}$.

C. Identification

We begin by discussing identification of the baseline price parameters ($\bar{\alpha}_{d(i)0}$), before discussing the other preference parameters.

We assume that the detailed fixed effects that we include in the model (in addition to the advertising controls) absorb taste variation that is relevant for price-setting. The brand-size effects, $\xi_{d(i)b(j)s(j)}^{(1)}$, absorb the influence of unobserved product attributes that are not captured by the included observable product attributes. Variation in taste for brands or particular sizes over time, due, for instance, to seasonal patterns, are captured by the brand-time, $\xi_{d(i)b(j)t}^{(2)}$, and size-time, $\xi_{d(i)s(j)t}^{(3)}$, effects. In addition, we control separately for product level advertising, which will capture the effect on demand of the (overwhelmingly national) advertising in the UK drinks market. Finally, tastes for brands or sizes may vary across retailers, which is captured by the brand-retailer, $\xi_{d(i)b(j)r}^{(4)}$, and size-retailer, $\xi_{d(i)s(j)r}^{(5)}$, effects.

The price variation that we exploit is *product-level time-series variation that is differential across retailers*. In particular, while differences in the average price (over time) that different retailers set for a given product are small, the degree of co-movement in prices for the same product in different retail chains over time is low. The differential price movements are generally driven by time-limited price reductions, which vary in timing and depth across retailers; see Appendix C. These price reduction strategies, which are agreed in advance at bilateral meetings between drinks firms and retailers (see Section II.D), creates, from consumers' perspective, randomness in the prices they face.

A threat to our identification strategy is that consumers respond to promotions by intertemporally switching their purchases, i.e., stocking up when the price is low and consuming from this stock when the price is high. This would lead to over-estimates of own price elasticities and likely to under-estimates of cross price elasticities (Hendel and Nevo (2006)). In Appendix C we show that, when purchasing on sale, consumers are more likely to choose a different brand, container type and size relative to their previous purchase, but they do not systematically change the timing of their purchases. In other words, sales in the UK drinks market primarily lead to *intra*-temporal substitution rather than intertemporally switching.

We do not model the decision that consumers make over which retailer to shop with, though we do allow brand and pack-size preferences to vary flexibly by

retailer. We therefore assume that the choice of retailer is not driven by time-varying, drink product specific factors. Drinks expenditure comprises only a small share of total grocery expenditure (4%), so we think this is a reasonable assumption. In addition, we find that when purchasing on sale, there is no economically meaningful increase in the likelihood that the consumers shopped with a different retailer compared with their previous purchase, which supports our assumption (see Appendix C).

For the non-price product attributes the baseline preference is identified as long as the attribute exhibits within brand-size variation (otherwise it is absorbed by the brand-size fixed effects). For instance, a given consumer group will have a stronger preference for sugar if, conditional on brand-size and other variables in the model, they more regularly choose sugary rather than diet products.

We observe in our micro data many consumers of different total dietary calories, household composition and income making repeated choice while facing different price vectors. We use this variation to incorporate rich preference heterogeneity into the model across these observable consumer attributes. We also include rich unobserved preference heterogeneity. The panel structure of our data is helpful in pinning down the spread parameters governing the unobserved preference distributions. For instance, conditional on the average price paid for chosen options across *all* choice occasion, the higher is the average *within-consumer* covariation in the price of chosen options across choice occasions, the higher is the dispersion in price preferences, which is captured by the spread parameter, $\sigma_{d(i)0}^\alpha$.

IV. Demand and supply estimates

A. Consumer substitution patterns

We estimate the demand model outlined in Section III.A using simulated maximum likelihood, and report the coefficient estimates in Appendix D. The estimated coefficients exhibit some intuitive patterns: those with relatively high overall added sugar in their diets have stronger preferences for sugary drinks products, and those with lower incomes are more sensitive to price, have stronger preferences for soft drinks and weaker preferences for pure fruit juice. The variance parameters of the random coefficients are significant both statistically and in size, indicating an important role for unobserved preference heterogeneity. We allow for correlation in preferences (within consumer group) between tastes for drinks (relative to the outside option) and preferences for sugar, and find a negative relationship.

The estimated preference parameters determine our demand model predictions of how consumers switch across products as prices change. The mean own-price elasticity is around -2.1 (in both the at-home and on-the-go segments), though with significant variation around this: 25% of products have own-price elasticities with magnitude greater than 2.5, a further 25% of products have own-price elasticities with magnitude less than 1.6. The distribution of the cross-price elasticities

exhibits a high degree of skewness, with the mean close to the 75th percentile. This reflects consumers' willingness to switch between products close together in product characteristic space.

To illustrate this, Table 1 shows product level elasticities associated with a price change for two popular sizes – a 2l bottle and a 8 pack of 330ml cans – of Coke Regular and Diet Coke. It shows the impact on demand for each of the 2l bottle and 8×330ml packs of Coke and Pepsi, and the mean elasticities for other (non-cola) sugar-sweetened and diet beverages, and for pure fruit juice. The table highlights a number of intuitive patterns: (i) consumers are more willing to switch across cola products of the same variety (e.g., within Regular) than they are to other varieties (e.g., Diet or Zero) or to non-cola drinks; (ii) consumers are more willing to switch between products of the same size/pack type than they are to different sizes; (iii) consumer substitution from sugary varieties of Coke to sugary non-cola drinks (both sugar-sweetened beverages and fruit juice) is stronger than it is from Diet Coke. In Appendix D we report further details of product level elasticities.

In Table 2 we summarize the effects of increasing the price of all sugar-sweetened beverages by 1%. This leads to a 1.41% fall in liters demanded of sugar-sweetened beverages. Around 33% of the reduction in demand for sugar-sweetened beverages is diverted to alternative drinks. As we discuss in Section I, if alternative products are supplied non-competitively the degree of switching to them in response to a marginal tax rise is an important determinant of the optimal policy. The diversion ratio in Table 2 does not directly tell us this since (i) it reflects only demand responses to a price change, but not supply side pricing responses to a tax change and (ii) it is evaluated at observed prices and not at the optimal tax rate. Nonetheless, as we show in Section V, the relatively high degree of substitution between the two product sets indicated by the diversion ratio plays an important role in determining optimal sugar-sweetened beverage taxation. The 1% increase in the price of sugar-sweetened beverages (at observed prices) leads to essentially no change in overall drinks expenditure.

B. Estimated costs and margins

We use the first order conditions of the firms' profit-maximization problem (equation (7)) to solve for product marginal costs, and hence the price-cost margins and Lerner indexes (margin over price) at observed prices. In Table 3 we show the averages of these for sugar-sweetened beverages and alternative products.

The average Lerner index is 0.57 for sugar-sweetened beverages and 0.55 for alternative products. This indicates that firms exercise a significant degree of market power when setting the prices both of sugar-sweetened beverages and alternative drinks. As we illustrate in Section V, failing to account for distortions from the exercise of this market power leads to substantial unrealized welfare gains when setting tax policy. In Appendix D we show that there is substantial variation in equilibrium margins across brand and that average margins are declining in

TABLE 1—SELECTED ELASTICITIES FOR COLA PRODUCTS

	Coke				Pepsi				Non-colas		
	Regular 2l b. 8 pk.	Diet 2l b. 8 pk.	Zero 2l b. 8 pk.	Max 2l b. 8 pk.	Regular 2l b. 8 pk.	Zero 2l b. 8 pk.	Max 2l b. 8 pk.	SSBs	Diet	Fruit juice	
<i>Regular</i>											
2l bottle	-2.204	0.018	0.011	0.009	0.023	0.017	0.012	0.009	0.007	0.003	
8x330ml can	0.036	-2.832	0.017	0.022	0.035	0.042	0.017	0.021	0.013	0.008	
<i>Diet</i>											
2l bottle	0.010	0.007	-2.185	0.014	0.010	0.007	0.020	0.014	0.003	0.003	
8x330ml can	0.014	0.018	0.026	-2.777	0.014	0.017	0.026	0.031	0.006	0.011	

Note: Numbers show price elasticities of market demand (for products listed in top row) in the most recent year covered by our data (2012) with respect to price changes for two specific pack sizes of Coke Regular and Diet Coke (shown in first column). “Non-colas” exclude Coke and Pepsi and are means over products belonging to each of the sets, sugar-sweetened beverages (SSBs), diet drinks and fruit juices.

TABLE 2—SWITCHING DUE TO AN INCREASE IN THE PRICE OF SUGAR-SWEETENED BEVERAGES

Own price elasticity for sugar sweetened beverages	-1.41 [-1.46, -1.37]
% lost demand diverted to substitute drinks	32.8 [32.3, 33.6]
% change in overall drinks expenditure	0.047 [0.033, 0.060]

Note: We simulate the effect of a 1% price increase for all sugar-sweetened beverage products. The first row shows the % reduction in volume demanded of sugar-sweetened beverages, the second row shows how much of the volume reduction is diverted to substitute drinks products, and the third row shows the percent change in total drinks expenditure. Numbers are for the most recent year covered by our data (2012). 95% confidence intervals are given in square brackets.

product size (since, on average, price per liter is strongly declining in size, while marginal cost per liter is flatter across the size distribution).

It is important to emphasize that distortions from the exercise of market power are endogenous to tax policy. For instance, a tax levied on sugar-sweetened beverages may exacerbate distortions from market power on these products if firms respond by raising their tax-exclusive prices (and hence product level margins) and/or consumers respond by downsizing to smaller high margin products.

TABLE 3—SUMMARY OF COSTS AND MARGINS

	Sugar-sweetened beverages	Alternative products
Price (£/l)	1.09	1.07
Marginal cost (£/l)	0.42	0.44
Price-cost margin (£/l)	0.67	0.63
Lerner index (margin/price)	0.57	0.55

Note: We recover marginal costs for each product in each market. The table shows the average price, marginal cost, price-cost margin (all expressed in per liter terms) and Lerner index among sugar-sweetened beverages and substitute products, constructed using quantity weights. We report the values for the most recent year covered by our data (2012).

C. Discussion of demand and supply estimates

Our demand and supply estimates enable us to capture rich consumer level substitution patterns across products, as well as product level price-cost margins. These are key inputs into our empirical study of optimal taxation with market power. However, this necessarily entails making identifying and functional form assumptions about the nature of demand and firm competition. Here we compare our estimates with those in the existing literature, and to alternative information

on price-cost margins. We also provide validation of our model using data on price changes following the introduction of the UK's soft drinks tax.

We estimate an own-price elasticity for sugar-sweetened beverages of 1.41. We calculate this by simulating an increase in the prices of all sugar-sweetened beverages by 1% and recovering the change in demand for those products, allowing for substitution between sugary beverages, to alternative drinks and to not buying drinks. Allcott, Lockwood and Taubinsky (2019a) employ an alternative approach, using US scanner data and an instrumental variable methodology applied to quarterly purchases of sugar-sweetened beverages. They estimate an own-price elasticity for sugar-sweetened drinks in line with ours (between -1.37 and -1.48, depending on the specification). The margins we recover imply the average Lerner index in the market is 0.56, which is broadly consistent with the gross margins reported in accounting data, which are between 35-70% (see Competition Commission (2013)).

PASS-THROUGH. — We use our demand and supply estimates to simulate how firms choose to adjust their margins when a tax is introduced. A potential concern with simulated pass-through of a hypothetical tax is that it can be influenced by functional form assumptions. In particular, the curvature of market level demand is an important determinant of tax pass-through (see Weyl and Fabinger (2013)). A feature of logit demand models with no heterogeneity in preference parameters is that they heavily restrict demand curvature. However, the addition of preference heterogeneity breaks the link between the curvature of individual- and market-level demand curves, allowing for more flexibility in the latter, as curvature now also depends on how the composition of individuals along the market demand curve changes (Griffith, Nesheim and O'Connell (2018)).

We provide direct evidence that our model succeeds in generating realistic tax pass-through predictions by simulating the introduction of the UK's Soft Drinks Industry Levy (SDIL) in 2018 and comparing this with what happened in practice – see Appendix D for full details. This tax was introduced more recently than the period covered by our data, so we use a weekly database of UPC level prices, collected from the websites of six major UK supermarkets, that cover the period 12 weeks before and 18 weeks after the introduction of the tax. We estimate the average within-product price changes before and after the introduction of the tax. Figure 3 shows the results for those products with sugar content above 8g per 100ml, which were subject to the higher rate (results for the low tax band and tax-exempt products are shown in Appendix D). We find evidence that the tax was slightly overshifted, with average pass-through rates of 105-108% and no change in the price of untaxed products. The price changes predicted by our model are very close to the observed price changes.

To illustrate the extent to which our pass-through predictions are driven by the logit shocks in our demand model we also simulate pass-through shutting down all preference heterogeneity in the model. Doing this results in simulated

tax pass-through that is around 50% on average. This illustrates that modeling preference heterogeneity is key in enabling us to capture realistic pass-through patterns.

Modest over-shifting of a drinks tax is broadly consistent with evidence from ex-post evaluations of these policies. For example, the Philadelphian tax was found to be fully passed through to prices (Seiler, Tuchman and Yao (2021), Cawley et al. (2020)), and in Mexico the tax was fully to slightly more than fully passed through to prices (Grogger (2017), Colchero et al. (2015)). An exception is Berkeley, where pass-through of the tax is estimated to be statistically insignificant or low (e.g. Rojas and Wang (2017), Bollinger and Sexton (2018)).

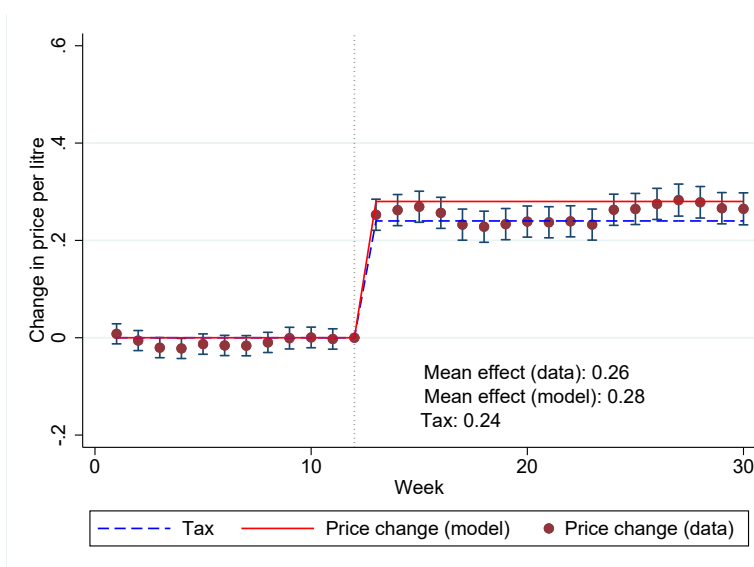


FIGURE 3. OUT-OF-SAMPLE MODEL VALIDATION: UK SOFT DRINKS INDUSTRY LEVY

Note: The red circles show the estimated price changes (relative to the week preceding the introduction of the tax). Appendix D contains full details of the specification. 95% confidence intervals shown. The blue line shows the value of the tax, and the red solid line shows the predicted price changes from our estimated demand and supply model.

V. Optimal tax results

In this section we combine our empirical model of the UK drinks market with the tax design framework that we outline in Section I to quantitatively solve for optimal tax policy; see Appendix E for the solution algorithm. We begin by considering the efficiency-maximizing sin tax rate, where the government seeks to minimize allocative distortions, but is indifferent to the incidence of the tax across individuals. Next we consider the optimal sin tax rate when the government has

distributional concerns, highlighting the importance of inequality in consumption and profit holdings across individuals in determining optimal policy. Finally, we consider the potential welfare gains from multi-rate taxation and levying tax directly on product sugar content rather than volumetrically.

EXTERNAL COSTS OF SUGAR-SWEETENED BEVERAGE CONSUMPTION. — The welfare effect of tax policy depends on the magnitude of externalities from sugar-sweetened beverages. Consumption of these products can increase health care costs, the bulk of which are not borne by the individual (see Scientific Advisory Committee on Nutrition (2015)). As in Allcott, Lockwood and Taubinsky (2019a), we measure the monetary value of the average externality using epidemiological evidence from Wang et al. (2012). They estimate the impact of a reduction in sugar-sweetened beverage consumption on US healthcare costs;²³ we translate this into a UK-equivalent figure. Provision of healthcare in the UK is almost entirely public, meaning that increased healthcare costs directly impact the government's budget. We also use evidence from the UK based on an epidemiological study (Briggs et al. (2013)) and information on the NHS costs of treating obesity. Both approaches provide a similar picture and lead to an estimate of the average externality of approximately 4 pence per 10g of sugar – see Appendix E for details of these calculations. Based on a review of the medical consequences of added sugar intake, the World Health Organization (2003) recommends that added sugar should make up less than 10% of dietary calories. We therefore assume that externalities from sugar-sweetened beverage consumption are due to those people with more than 10% of their dietary calories from sugar; this implies the externality per 10g of sugar consumption is 5 pence for the 80% of consumers with dietary sugar above 10% of their calorie intake. The marginal externality from sugar-sweetened beverage consumption therefore varies across products based on their sugar content, and varies across consumers based on the total amount of added sugar in their diet. The average marginal externality due to 10 oz of sugar-sweetened beverage consumption, for those with dietary sugar exceeding 10% of total calories, is 14 pence. Below we show how our results vary with the nature of the mapping between sugar consumption and externalities.

If a marginal change in drinks tax policy causes consumption changes outside the drinks market, distortions from the exercise of market power in the supply of non-drinks products will impact optimal policy. In our baseline results we assume that the numeraire good is competitively supplied. However we show below that, since the impact of a marginal tax change on total drinks spending (and hence numeraire good consumption) is small, our results are numerically insensitive to this assumption.

²³The majority of these savings come from the reduced costs of treating diabetes and cardiovascular disease. They do not account for the potential impact on, for example, social security from people living longer.

A. Efficiency-maximizing policy

Efficiency-maximizing tax policy minimizes the allocative distortions resulting from governmental (distortionary tax) and non-governmental (externality and market power) distortions. It is indifferent to the distribution of welfare gains across individuals. The efficiency-maximizing policy corresponds to the maximum of the social welfare function (equation (1)) when all individuals have social marginal welfare weights equal to one.

To illustrate the impact of market power on the efficiency properties of sin taxation, we consider a single tax rate applied to the set of sugar-sweetened beverages (i.e., the sin products, which comprise product set \mathcal{S}). We summarize the impact of the tax in the first row of Table 4. The efficiency-maximizing tax rate is 4.19 pence per 10 oz; at the time of writing, US and UK sugar-sweetened beverage taxes range from 7 to 15 pence per 10 oz. The tax leads to an average increase in the price of sugar-sweetened drinks of 19.0%, and little change, on average, in the price of alternative drinks. This results in a 28.7% reduction in consumption of sugar-sweetened beverages, and a 7.0% increase in the consumption of alternative drinks.²⁴ The tax leads to substantial losses in consumer surplus (£510m, compared to total expenditure in the overall drinks market of approximately £9b) and moderate profit losses (£190m, or around 3.5% fall in market variable profits). However, this is more than made up for by a £386m (25%) fall in externality costs and £409m in excise tax revenue. Overall, economic efficiency increases by £94m.

The efficiency-maximizing tax rate consists of two components – one reflecting distortions from externalities, the other distortions from the exercise of market power for sin and alternative (substitute) products. To illustrate the role that each plays in determining efficiency-maximizing tax policy, it is instructive to consider two naive policies, one that ignores distortions from market power for all goods, and one that ignores them for alternative products.²⁵

A government that completely ignores distortions from market power would set a tax rate equal to 11.97 pence per 10 oz. This policy ignores the fact that equilibrium prices are set in excess of marginal costs and results in a sub-optimally high rate. The second row in Table 4 shows that this leads to a fall in consumer surplus and profits that is over twice as large as under efficiency-maximizing policy, with the combined loss outweighing the fall in externality costs plus tax revenue, meaning that overall economic efficiency falls by £71m. A government that takes account of market power distortions but only among the taxed sin products, would choose to set a tax rate that is approximately zero; the positive

²⁴This fall in sugar-sweetened beverage consumption is due to an 14.3% reduction in the probability, on average, a consumer purchases from this product set and, conditional on buying, a reduction in volume of 11.9%.

²⁵In the first case the government sets τ that satisfies $\tau = \bar{\phi} + \frac{n(\mathcal{S})}{dX_{\mathcal{S}}/d\tau_{\mathcal{S}}} \text{cov}(\phi_j^i, dx_j^i/d\tau_{\mathcal{S}})$, in the second case it sets the rate that satisfies $\tau = \bar{\phi} + \frac{n(\mathcal{S})}{dX_{\mathcal{S}}/d\tau_{\mathcal{S}}} \text{cov}(\phi_j^i, dx_j^i/d\tau_{\mathcal{S}}) - \bar{\mu}_{\mathcal{S}}$.

equilibrium margins for the sin goods offset the externality associated with their consumption. However, this policy results in a sub-optimally low tax rate. It ignores the fact that other products (substitutes to sin products) are supplied non-competitively. As we show in Tables 2 and 3, substitution to alternative drinks products is substantial, and these products (like sin products) have substantial mark-ups. Efficiency-maximizing policy depends on the average margins for sin goods relative to those on alternatives.

B. Optimal policy with distributional concerns

When the government has distributional concerns, it must balance efficiency with equity considerations. In this case, the distribution of the effects of tax policy across individuals, which depends both on the distribution of consumption and profit holdings, matters for optimal policy.

Profits flow to the government (via corporate and dividend taxes), and to domestic and overseas residents. Measuring stock ownership across the income distribution is challenging and remains a topic of considerable debate. Recent papers, including Saez and Zucman (2016) and Smith, Zidar and Zwick (2020), use a combination of dividend income and realized capital gains to estimate wealth in publicly traded stocks. In this spirit, we use information from the UK national accounts and the distribution of dividend income to allocate profits to different groups. The corporate (see Bilicka and Devereux (2012)) and effective average dividend tax rates leads to the government collecting 29% of profits. Using data from the national accounts, we set the fraction flowing overseas to 30%. We assume that the remaining 41% is distributed to UK residents in proportion to the share of (net-of-tax) dividend income received by households in equivalized income bands. The net-of-tax profit holdings of domestic residents is concentrated among the relatively wealthy; households with equivalized income below £10k make up around 25% of the population, but receive less than 3% of post-tax domestic dividend income; households with equivalized income above £45k comprise 5% of the population, but receive more than 21%. Sugar-sweetened beverage taxation will mainly affect the profits of drinks firms. We assume that profit holdings in these firms are approximated by profit holdings in the economy more generally (which is reasonable based on diversified investment portfolios). However, as a robustness exercise we also show results in the case when post-tax profits flow to individuals with zero marginal social welfare weights. In Appendix E, we provide full details of these calculations.

We parameterize the social marginal welfare weights as $g^i = (\tilde{y}_i)^{-\vartheta}$, where \tilde{y}_i is equivalized household income and ϑ captures the degree of inequality aversion in government preferences. As our baseline we set $\vartheta = 1$. In all calculations, we assume the government places a social marginal welfare weight of zero on the portion of profits that flows to overseas individuals.

We first illustrate how the distribution of consumption and profit holdings affect

TABLE 4—EFFICIENCY-MAXIMIZING SINGLE RATE SIN TAX POLICY

	Tax rate (p/10oz)	% change in		Change (relative to zero tax) in: (£m)				
		Price	Cons.	Consumer surplus	Total profits	Excise tax revenue	External cost savings	Total efficiency
Optimal	4.19 [4.02, 4.54]	19.0% [19.3%, 18.4%]	-28.7% [-27.8%, -29.3%]	-510 [-526, -494]	-190 [-198, -181]	409 [398, 420]	386 [361, 402]	94 [77, 110]
Pigouvian	11.97 [11.81, 12.17]	55.8% [58.2%, 54.4%]	-59.2% [-58.2%, -60.2%]	-1199 [-1244, -1167]	-429 [-446, -407]	749 [733, 772]	808 [765, 842]	-71 [-104, -39]

Note: Optimal refers to efficiency-maximizing policy. Pigouvian refers to policy set by a government that ignores distortions from the exercise of market power. Price and consumption changes are for sugar-sweetened beverages. Consumer surplus, Total profits, External costs, Excise tax revenue and Total efficiency numbers are per annum and report the change relative to no drinks taxation. Total profits are inclusive tax revenues from taxation of corporate profits. Total efficiency = Consumer surplus+Total profits+Excise tax revenue+External cost savings. 95% confidence intervals are given in square brackets.

policy for a fixed set of government preferences for equity, before showing how the strength of these preferences affect the optimal rate. In the final part of this section, we show how the market structure and nature of externalities affects optimal policy.

HOW THE DISTRIBUTION OF CONSUMPTION AND PROFITS AFFECTS POLICY. — In Table 5 we summarize the impact of distributional concerns on the optimal sin tax rate. The first row of the table shows the optimal rate, and its impact on welfare, under our measure of the true distribution of profit holdings (as described above). The optimal tax rate in this case is 5.97 pence per 10 oz – over 40% higher than the efficiency-maximizing rate (4.19p/10oz) – and it achieves an increase in social welfare of £167m. There are three channels through which distributional concerns affect the optimal rate. We conduct two thought experiments, based on counterfactual distributions of profits, to highlight the relative importance of these different channels.

First, we consider what the optimal rate would be if the government were to collect all profits as tax revenue (row two of the table). This isolates the impact of distributional concerns arising purely from consumption patterns. Since sugar-sweetened beverages are more popular with low income households, the optimal tax in this case lies *below* the efficiency-maximizing rate (it is 3.23 pence per 10 oz). A similar effect is highlighted by Allcott, Lockwood and Taubinsky (2019*a*). Second, we consider optimal policy if the share of profits flowing to overseas individuals equals the true share, but the government collects all domestically owned profits as tax revenue (row three). This introduces a second channel through which distributional concerns impact optimal policy; since a share of profits are owned by foreigners, who, under our social preference specification, are assigned social marginal welfare weights of zero, the government places less weight on the fall in profits resulting from tax policy. This leads the optimal tax rate to increase from 3.23 (when government is assumed to tax fully foreign as well as domestic profits) to 4.71 pence per 10 oz.

Comparing the final row with the first illustrates the importance of the third channel – the impact of the unequal distribution of domestic profits in combination with the government’s preference for equity. This leads the optimal tax rate to increase further (from 4.71 to 5.97 pence per 10 oz). Domestic profits are disproportionately in the hands of high income households, meaning the incidence of profit losses associated with the sin tax is mainly on those with relatively low social marginal welfare weights. All else equal, this increases the progressivity of the tax, raising the optimal rate and the size of associated welfare gains.

In summary, relative to an efficiency-maximizing government, distributional concerns lead to an increase in the optimal tax rate. This is because post-tax profits are mainly in the hands of foreign and relatively wealthy domestic residents, who have relatively low or zero social marginal welfare weights. This effect more than offsets the influence of distributional concerns over consumption pat-

TABLE 5—IMPACT OF DISTRIBUTIONAL CONCERNS ON OPTIMAL SINGLE RATE SIN TAX POLICY

	Tax rate (p/100z)	Welfare components (£m)					Total welfare
		Private welfare, from:	Tax revenue:	Ext. cost savings			
		Cons.	Profits	Sin tax	Profit tax		
True profit distribution	5.97 [5.51, 6.37]	-747 [-770, -726]	-43 [-45, -41]	522 [510, 537]	-74 [-77, -71]	509 [477, 529]	167 [142, 188]
All profits taxed at 100%	3.23	-439	0	336	-152	311	56
Domestic profits taxed at 100%	2.84, 3.56 4.71 [4.28, 5.05]	[-453, -427] -611 [-630, -595]	- 0 -	[326, 344] 445 [434, 457]	[-158, -144] -147 [-153, -140]	[291, 325] 424 [397, 441]	[42, 69] 111 [90, 128]

Note: Numbers summarize the effect of policy when the social marginal welfare weight on foreign individuals is 0 and on domestic individuals is $1/g$, showing effects under the true distribution of profit holdings (row (1)) and counterfactual distributions (rows (2) and (3)). Welfare numbers are per annum and report the change relative to no drinks taxation. Total welfare = Private welfare+Tax revenue+External cost savings. 95% confidence intervals are given in square brackets.

terns, which, all else equal, act to lower the optimal rate since sugar-sweetened beverages are disproportionately consumed by relatively low income (high social marginal welfare weight) consumers.

HOW THE STRENGTH OF PREFERENCES FOR EQUITY AFFECTS POLICY. — In Figure 4 we summarize how differences in the strength of social preferences for equity impact the optimal tax rate (panel (a)), the associated welfare gain (panel (b)), and the fraction of possible welfare gains forgone if the government ignores distortions from the exercise of market power when setting policy²⁶ (panel (c)). On the horizontal axis we plot efficiency-maximizing policy, and policy when the social marginal welfare weight on foreign profits is 0 and the weights on domestic consumers are $(\tilde{y}_i)^{-\vartheta}$ for $\vartheta = \{0, 1, 2, 3, 4\}$.

Relative to the efficiency-maximizing policy, when $\vartheta = 0$, the government places a social marginal welfare weight of zero on foreign profits, which acts to raise the optimal tax rate. When $\vartheta > 0$, the government is no longer indifferent to the incidence of the tax across domestic individuals. Thus, the fact that consumption is concentrated among those with low incomes, all else equal, acts to lower the optimal rate. However, this is offset by the concentration of post-tax domestic profit holdings among those with high incomes (meaning they bear a disproportionately large share of profit reductions from sin taxation), which acts to raise the optimal rate. ϑ controls the strength these effects exert on optimal policy. The optimal rate (and associated welfare gain) peaks at $\vartheta = 1$. When $\vartheta > 1$ the impact of consumption inequality on the optimal rate dominates the effect of inequality in domestic profits holding. Overall, the optimal tax rate and associated welfare gain are relatively insensitive to the inequality aversion parameter, and for all values of ϑ , the optimal tax rate and associated welfare gain are larger than under efficiency-maximizing policy.

Panel (c) shows that the welfare losses from ignoring distortions due to the exercise of market power are mitigated under social preferences for equity, but they remain substantial. Policy set by an efficiency-maximizing planner that ignores market power distortions leads to forgone welfare gains of more than 150%. The forgone welfare gains from this form of naive policy when there are distributional concerns are between 64% (when $\vartheta = 0$) and 35% (when $\vartheta \geq 2$). With stronger preferences for equity, since post-tax profits are disproportionately in the hands of those with relatively low or zero social marginal welfare weights, the government's welfare function places less weight on profits (than under efficiency-maximization), and therefore ignoring market power is less costly. However, even with strongly inequality averse preferences there remains a substantial loss from ignoring market power — this is because some profits are collected by the government as tax revenue, and because low and moderate income consumers have small, but non-zero, profit holdings.

²⁶Solving for the τ that satisfies $\tau = \bar{\phi} + \frac{n(S)}{d\bar{x}_S/d\tau_S} \text{cov}(\phi_j^i, dx_j^i/d\tau_S) + \frac{1}{d\bar{x}_S/d\tau} \text{Cov}(g^i, \sum_{j \in S} x_j^i)$.

Another possibility is that, while the government has inequality averse preferences with respect to consumption, it places a social marginal welfare weight of zero on all *post*-tax profits, perhaps due to uncertainty in measuring the distribution of post-tax profits across individuals. Even in this case, market power and the realization of profits has an important bearing on optimal sin tax policy, because sin taxation leads to spillovers to the profit (i.e., corporate and dividend) tax bases. The optimal sin tax rate in this case is 6.90 pence per 10 oz. If the government ignores market power when setting the tax rate (and hence the spillovers to the corporate and dividend tax bases) it would set a rate of 11.26, which would result in unrealized welfare gains of 26% (see row (2) in Table 6). This illustrates that, even if post-tax profits holdings are more concentrated in the hands of the rich than our measure of the profit distribution suggests (which is based on post-tax dividend income), market power continues to have an important bearing on optimal tax policy.

HOW MARKET STRUCTURE AFFECTS POLICY. — Our baseline results assume that the numeraire good is competitively supplied. However, it is possible that when consumers switch away from drinks they switch to other non-competitively supplied goods. This acts to raise the optimal tax rate through an efficiency channel, as it dampens the market power correction element of the optimal tax (which equals the average margin of sin goods relative to alternatives). Yet it acts to lower the optimal rate through an equity channel, as switching to a non-competitively supplied numeraire good leads to off-setting profit gains mainly for the rich. Row (3) in Table 6 shows results when the numeraire good is supplied non-competitively, with a margin equal to that implied by the estimate of the UK economy-wide mark-up in De Loecker and Eeckhout (2018). It shows that the optimal tax rate is 6% higher and the associated welfare gain marginally higher, than in our baseline with a competitively supplied numeraire. The modest impact of adding market power for the numeraire on optimal policy is due to the relatively small fall in overall drinks expenditure (and hence rise in numeraire good consumption) induced by a marginal tax rise.

In row (4) of Table 6 we consider optimal tax policy under the counterfactual market structure of a perfectly competitive drinks market. Without any tax in place, moving from the true market structure to perfect competition raises welfare by £622m; the gains from eliminating market power distortions dwarf the resulting increase in externality costs. The optimal tax rate under perfect competition is 10.18 pence per 10 oz and the resulting welfare gain is £515 million, which is around 3 times as large as under the true market structure (row (1)). Competition and optimal tax policy therefore exhibit a form of complementarity.

HOW THE NATURE OF EXTERNALITIES AFFECTS POLICY. — Rows (5)-(10) of Table 6 show how changes in the shape of the function mapping sugar-sweetened beverage consumption to externalities impact tax policy. Rows (5) and (6) vary the overall

magnitude of external costs, setting them at 25% above and 25% below the baseline value, implied by Wang et al. (2012). Higher consumption externalities lead to an increase in the optimal tax rate and the associated welfare gains. However, even with larger externality costs, the losses from ignoring distortions from the exercise of market power remain substantial (at 23%).

Row (7) and (8) show results when we vary the convexity of the externality function (holding fixed the average marginal externality across consumers). Row (7) assumes marginal externalities are equal for all consumers; row (8) assumes that only those with total dietary sugar above 15% (compared with 10% in the baseline) generate externalities through their consumption. The consumption of people with higher overall dietary sugar is moderately more responsive to tax changes. Therefore, the more concentrated externality generation is among overall high sugar consumers, the more effective the sin tax is at reducing externality distortions. Hence, more convexity in the externality function leads to a higher optimal rate and larger resulting welfare gains.

We also consider the consequence of externality spillovers. In particular, consumers respond to the sugar-sweetened beverage tax by, in part, switching to pure fruit juices and flavored milk. These are typically exempt from drinks taxes partly because these products contain other (positive) nutrients that may offset negative consequences of sugar intake. However, this is subject to debate. In row (9) we show the consequences for optimal policy when externalities are associated with the untaxed alternative sources of sugar. This type of externality leakage diminishes the effectiveness of tax in reducing externality distortions, leading to a lower optimal rate and smaller welfare gains (relative to the baseline, where consumption of these alternative products is assumed not to generate externalities). Row (10) quantifies the gains of including pure fruit juices and flavored milk in the tax base, in the case where consumption of these products is associated with externalities. This results in a higher optimal rate and welfare gains that are £43m larger than when these products generate externalities but are not part of the tax base (row (9)).

In the final row of Table 6 we illustrate how optimal policy changes if social costs also include internalities. We assume externalities are the same as in the baseline case, but that there are also internalities, such that the average social cost of 10g of additional sugar consumption is 25% larger than in the baseline (and the same as in row (5)). We assume that the internality associated with an additional unit of sugar consumption (in monetary terms) for someone at the bottom of the income distribution is 30% larger than for someone at the top (based on the evidence from Allcott, Lockwood and Taubinsky (2019a)). In this case, the optimal tax rate is 8.85 pence per 10 oz. This is marginally higher than the optimal rate when the additional 25% of social costs are due to externalities (row (5)). This is because the presence of internalities (negatively related to income in severity) act to make the tax more progressive, raising the optimal rate. It remains the case that ignoring market power results in substantial forgone welfare gains.

TABLE 6—OPTIMAL TAX POLICY AND THE COSTS OF IGNORING MARKET POWER

	Optimal policy		Ignoring market power	
	Tax rate (p/10oz)	Change in welfare (£m)	Tax rate (p/10oz)	Change in welfare (£m)
(1) Baseline	5.97 [5.51, 6.37]	167 [142, 188]	11.26	100 [58, 129]
(2) Zero weight on post-tax profits	6.90 [6.42, 7.36]	213 [185, 238]	11.26	157 [114, 188]
Market structure				
(3) Numeraire good market power	6.34 [5.76, 6.87]	169 [140, 194]	11.26	116 [70, 150]
(4) Perfect competition	10.18 [9.73, 10.67]	515 [457, 567]	10.18	515 [457, 567]
Externalities				
(5) 25% larger	8.44 [7.87, 8.84]	312 [274, 343]	14.63	239 [187, 281]
(6) 25% smaller	3.49 [3.05, 3.81]	61 [47, 72]	8.01	-2 [-31, 18]
(7) Linear	5.35 [4.92, 5.73]	138 [116, 159]	10.33	73 [37, 105]
(8) More convex	7.88 [4.33, 5.48]	267 [226, 314]	13.98	194 [139, 262]
(9) Leakage	5.24 [4.78, 5.63]	153 [131, 172]	10.12	105 [69, 132]
(10) Broader base (+leakage)	6.25 [5.77, 6.77]	196 [172, 223]	12.97	109 [66, 152]
(11) +Internalities	8.85 [8.35, 9.22]	338 [303, 367]	15.23	263 [217, 303]
				22% [17, 28]

Note: The first two columns summarize optimal policy, the final three summarize policy set by government that ignores distortions from market power (with the final column showing the % of welfare gains from optimal policy forgone). All numbers are based on social marginal welfare weights on foreign individuals of 0 and on domestic individuals of $1/\bar{g}$. Row (1) repeats numbers under our central calibration of the numeraire good margin and externalities. Row (2) sets social marginal welfare weights on post-tax profits to zero. The remaining rows present results under alternative market structures and externality functions, with details described in the text. Welfare numbers are per annum and report the change relative to no drinks taxation. 95% confidence intervals are given in square brackets.

C. *Optimal policy with alternative tax instruments*

Our focus to this point is on a volumetric single tax rate levied on sugar-sweetened beverages. This form of taxation is very common among the taxes that have actually been implemented.²⁷ In this section we consider alternative tax instruments. In particular, we consider a system of multiple volumetric tax rates applied in the drinks market, and a tax that is levied directly on product sugar content.

MULTI-RATE SYSTEM. — Alcohol taxation typically involves separate rates for different alcohol types. We consider a similar system applied in the (non-alcoholic) drinks market, in which the government can set rates that vary across broad drink types (for instance colas, lemonades etc.); see Appendix E for details. Under the constraint that policy cannot lead to a deterioration in the government’s budget, inclusive of the budgetary externality, optimal policy involves subsidizing non-sugar sweetened drinks paid for by a combination of taxes on sugar-sweetened drinks and reductions in externalities. Overall it raises welfare by 80% more than than under the optimal single rate system. When subsidies are prohibited, optimal policy entails varying the tax rate across different sugar sweetened beverage types. Relative to under the optimal single sugar sweetened beverage tax rate, the average price rise for sin products and fall in externalities is similar. However, rate differentiation enables the multi-rate system to better target inefficiencies, resulting in a smaller loss in private welfare than under the single rate system and a welfare rise that is 17% higher.

A SUGAR TAX. — An alternative to the volumetric sugar-sweetened beverage taxes typically used is to levy a tax directly on the sugar in drinks. This has the advantage that the tax is proportional to the total quantity of the externality generating attribute that a product contains. We compute the optimal tax on the sugar in sweetened beverages, assuming firms do not change their products’ sugar contents. Compared with the optimal single rate volumetric tax, it leads to slightly higher prices, larger falls in consumer welfare and lower tax revenue, however this is more than offset by larger reductions in externalities. Overall, the sugar tax raises welfare by 23% more than the single rate volumetric tax (see Table E.4 in Appendix E). The forgone welfare gains from ignoring market power are similar under the sugar and volumetric taxes.

Another possible advantage of a sugar tax is that it would incentivize input substitution, as a firm can lower its exposure to the tax by reducing the sugar content of its products. The impact of input substitution on welfare will depend on how

²⁷As of mid 2021, of the 44 countries and 9 US cities that have implemented sugar-sweetened beverage taxes, only Mauritius, South Africa and Sri Lanka have taxes levied directly on sugar content. In the large majority of cases, sugar-sweetened beverage taxes entail a single rate. Exceptions are the UK and Portugal, which have banded systems with two rates.

the change impacts production costs and consumers' valuation of the products. To illustrate this we follow Barahona et al. (2020) by assuming firms can lower sugar content in their products at the expense of raising production costs, but without altering product taste (full details are in Appendix E). Our simulations suggest that the smaller is the increase in production costs from removing sugar, the larger are the welfare gains associated with a sugar tax over a volumetric tax rise. While firms choose to reformulate products to maximize profits, the associated externality reductions are sufficiently high that social welfare increases. An important avenue for future research is to estimate the consumer welfare impact of firms altering their products to avoid taxes.²⁸

VI. Conclusion

In this paper we show how market power affects the efficiency and redistributive properties of sin taxation. Allocative distortions from the exercise of market power lead optimal policy to depend on the magnitude of equilibrium price-cost margins on sin products, relative to alternatives. The relative concentration of profit holdings in the hands of the wealthy, leads policy to be more progressive than if no profits were realized, counteracting the regressive incidence of the tax based on consumption patterns alone. We show, in an application to sugar sweetened beverage taxation, that market power exerts a quantitatively significant impact on optimal policy and ignoring it leads to substantial unrealized welfare gains. We believe our results are both of direct relevance for the design and implementation of sugar-sweetened beverage taxation and yield important lessons for the design of sin taxes levied in other specific markets. We conclude by suggesting two promising avenues for future research.

One rationale for sin taxation is to correct for consumer misoptimization. When firms have market power they may seek to exploit consumer mistakes, for instance, through persuasive advertising, exploiting consumer self-control problems or obfuscating the unhealthy nature of products (e.g., see Spiegler (2010)). An interesting avenue for future research is to explore the interaction between consumer misoptimization and firms' strategies, and the implications for sin tax design.

We focus in this paper on illustrating how the non-competitive structure of a market can have an important bearing on the welfare effects of taxes levied in that market. We show how optimal policy varies with competitive conditions, as well as for different levels of profit taxes, under the assumption that the government optimizes sin taxation while treating competition policy fixed. When considering tax policy in a single, relatively small, market this assumption is both reasonable and realistic. However, in other contexts, when tax policy directly affects a large swath of the economy, such as that targeted at climate change, the case for

²⁸A hand full of countries have adopted hybrid taxes, which are volumetric, but with rates that vary across a small number of bands defined on the basis of product sugar content. Another interesting avenue for future research is to study whether these hybrid systems outperform a sugar tax. The recently introduced Soft Drink Industry Levy in the UK is one possible setting in which to study this question.

(and gains from) also adjusting competition policy are likely to be larger. Several papers highlight the importance of market power in markets associated with greenhouse gas emissions (e.g, see Bushnell, Mansur and Saravia (2008) for electricity and Hastings (2004) for gasoline). A second promising avenue for future work is to consider the joint determination of commodity taxation with other aspects of policy, including competition policy and corporate taxation, in tackling systemic externalities, such as climate change.

REFERENCES

- AC Nielsen.** 2007-2012. “AC Nielsen Advertising Digest.” <https://www.nielsen.com/solutions/media-planning/ad-intelligence/>, (accessed August 8, 2013).
- Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky.** 2019a. “Regressive Sin Taxes, with an Application to the Optimal Soda Tax.” *Quarterly Journal of Economics*, 134(3): 1557–1626.
- Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky.** 2019b. “Should We Tax Soda? An Overview of Theory and Evidence.” *Journal of Economic Perspectives*, 33(3): 23.
- Auerbach, Alan J., and James R. Hines.** 2002. “Taxation and Economic Efficiency.” In *Handbook of Public Economics*. Vol. 3, 1347–1421.
- Barahona, Nano, Cristobal Otero, Sebastian Otero, and Joshua Kim.** 2020. “Equilibrium Effects of Food Labeling Policies.” *Available at SSRN: 3698473*.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. “Automobile Prices in Market Equilibrium.” *Econometrica*, 63(4): 841–890.
- Berry, Steven T.** 1994. “Estimating Discrete-Choice Models of Product Differentiation.” *RAND Journal of Economics*, 25(2): 242–262.
- Bilicka, Katarzyna, and Michael Devereux.** 2012. “CBT Corporate Tax Ranking 2012.” *Oxford University Centre for Business Taxation*, 37.
- Bollinger, Bryan K., and Steven Sexton.** 2018. “Local Excise Taxes, Sticky Prices, and Spillovers: Evidence from Berkeley’s Soda Tax.” *Available at SSRN: 3087966*.
- Bonnet, Céline, and Pierre Dubois.** 2010. “Inference on Vertical Contracts between Manufacturers and Retailers Allowing for Nonlinear Pricing and Resale Price Maintenance.” *RAND Journal of Economics*, 41(1): 139–164.
- Bonnet, Céline, and Vincent Réquillart.** 2013. “Tax Incidence with Strategic Firms in the Soft Drink Market.” *Journal of Public Economics*, 106: 77–88.

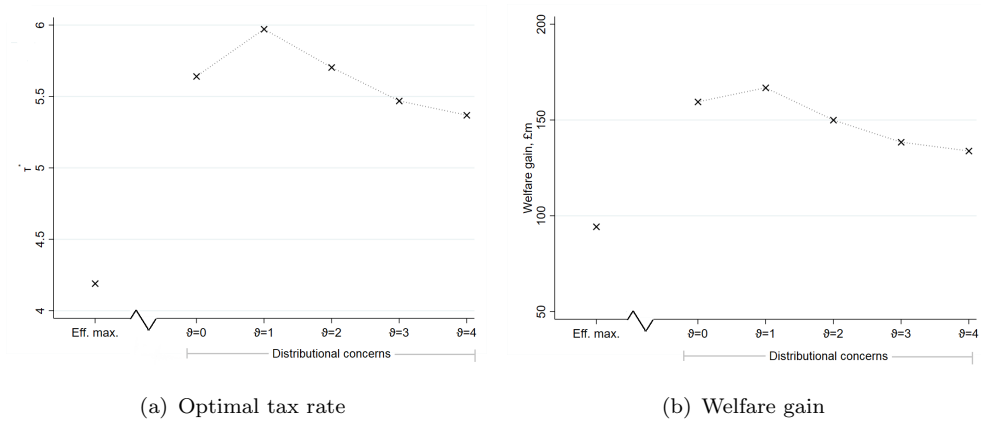
- Bresnahan, Timothy F.** 1989. "Empirical Studies of Industries with Market Power." In *Handbook of Industrial Organization*. Vol. 2, 1011–1057.
- Briggs, Adam D. M., Oliver T. Mytton, Ariane Kehlbacher, Richard Tiffin, Mike Rayner, and Peter Scarborough.** 2013. "Overall and Income Specific Effect on Prevalence of Overweight and Obesity of 20% Sugar Sweetened Drink Tax in UK: Econometric and Comparative Risk Assessment Modelling Study." *BMJ*, 347: f6189.
- Buchanan, James M.** 1969. "External Diseconomies, Corrective Taxes, and Market Structure." *American Economic Review*, 59(1): 174–177.
- Bushnell, James B., Erin T. Mansur, and Celeste Saravia.** 2008. "Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets." *American Economic Review*, 98(1): 237–266.
- Cawley, John, David Frisvold, Anna Hill, and David Jones.** 2020. "The Impact of the Philadelphia Beverage Tax on Prices and Product Availability." *Journal of Policy Analysis and Management*, 39(3): 605–628.
- Chernozhukov, Victor, Jerry Hausman, and Whitney Newey.** 2019. "Demand Analysis with Many Prices." *NBER Working Paper 02138*.
- Chetty, Raj.** 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics*, 1: 451–488.
- Colchero, M. Arantxa, Juan Carlos Salgado, Mishel Unar-Munguía, Mariana Molina, Shuwen Ng, and Juan Angel Rivera-Dommarco.** 2015. "Changes in Prices After an Excise Tax to Sweetened Sugar Beverages Was Implemented in Mexico: Evidence from Urban Areas." *PLOS ONE*, 10(12): e0144408.
- Competition Commission.** 2000. *Supermarkets: A Report on the Supply of Groceries from Multiple Stores in the United Kingdom*. London: The Stationary Office.
- Competition Commission.** 2013. *AG Barr / Britvic Merger Inquiry*. London: The Stationary Office.
- Conlon, Christopher T, and Nirupama S Rao.** 2015. "The Price of Liquor Is Too Damn High: Alcohol Taxation and Market Structure." *NYU Wagner Research Paper No. 2610118*.
- Cooper, Michael, John McClelland, James Pearce, Richard Prisinzano, Joseph Sullivan, Danny Yagan, Owen Zidar, and Eric Zwick.** 2016. "Business in the United States: Who Owns It and How Much Tax Do They Pay?" *Tax Policy and the Economy*, 41.

- De Loecker, Jan, and Jan Eeckhout.** 2018. "Global Market Power." *NBER Working Paper 24768*.
- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger.** 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics*, 135(2): 561–644.
- Diamond, Peter A.** 1973. "Consumption Externalities and Corrective Imperfect Pricing." *The Bell Journal of Economics*, 4(2): 526–538.
- Dubois, Pierre, Rachel Griffith, and Martin O'Connell.** 2018. "The Effects of Banning Advertising in Junk Food Markets." *Review of Economic Studies*, 85(1): 396–436.
- Dubois, Pierre, Rachel Griffith, and Martin O'Connell.** 2020. "How Well Targeted Are Soda Taxes?" *American Economic Review*, 110(11).
- Einav, Liran, Amy Finkelstein, and Pietro Tebaldi.** 2019. "Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies." *Mimeo*.
- Einav, Liran, and Jonathan Levin.** 2010. "Empirical Industrial Organization: A Progress Report." *Journal of Economic Perspectives*, 24(2): 145–162.
- Fowlie, Meredith, Mar Reguant, and Stephen P Ryan.** 2016. "Market-Based Emissions Regulation and Industry Dynamics." *Journal of Political Economy*, 124(1): 55.
- Gorman, W. M.** 1980. "A Possible Procedure for Analysing Quality Differentials in the Egg Market." *Review of Economic Studies*, 47(5): 843–856.
- Griffith, Rachel, Lars Nesheim, and Martin O'Connell.** 2018. "Income Effects and the Welfare Consequences of Tax in Differentiated Product Oligopoly." *Quantitative Economics*, 9(1): 305–341.
- Griffith, Rachel, Martin O'Connell, Kate Smith, and Rebekah Stroud.** 2019. "The Evidence on the Effects of Soft Drinks Taxes." *IFS Briefing Note 255*.
- Grogger, Jeffrey.** 2017. "Soda Taxes And The Prices of Sodas And Other Drinks: Evidence From Mexico." *American Journal of Agricultural Economics*, 99(2): 481–498.
- Harding, Matthew, and Michael Lovenheim.** 2017. "The Effect of Prices on Nutrition: Comparing the Impact of Product- and Nutrient-Specific Taxes." *Journal of Health Economics*, 53: 53–71.
- Hastings, Justine S.** 2004. "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California." *American Economic Review*, 94(1): 317–328.

- Hendel, Igal, and Aviv Nevo.** 2006. “Measuring the Implications of Sales and Consumer Inventory Behavior.” *Econometrica*, 74(6): 1637–1673.
- Hendren, Nathaniel.** 2016. “The Policy Elasticity.” *Tax Policy and the Economy*, 30(1): 51–89.
- Joint WHO/FAO Expert Consultation,** ed. 2003. *Diet, Nutrition, and the Prevention of Chronic Diseases. World Health Organisation Technical Report Series*, Geneva:World Health Organization.
- Kantar.** 2008-2012. “Kantar Worldpanel and Food On-The-Go Survey.” <https://www.kantar.com/expertise/consumer-shopper-retail/consumer-panels/consumer-goods-panel> and <https://www.kantar.com/expertise/consumer-shopper-retail/consumer-panels/out-of-home-panel?par=global/Consumer-Panels/Out-of-Home>, (accessed January 9, 2018).
- Kaplow, Louis.** 2021. “Market Power and Income Taxation.” *American Economic Journal: Economic Policy*, 13(4): 329–354.
- Kleven, Henrik J.** 2021. “Sufficient Statistics Revisited.” *Annual Review of Economics*, 13: 33.
- Lancaster, Kelvin.** 1971. *Consumer Demand: A New Approach*. Columbia University Press New York.
- Levell, Peter, Martin O’Connell, and Kate Smith.** 2016. “Chapter 9: Excise Duties.” In *Green Budget 2016*. Vol. 2016, 201–231. London:IFS.
- Luco, Fernando, and Guillermo Marshall.** 2020. “The Competitive Impact of Vertical Integration by Multiproduct Firms.” *American Economic Review*, 110(7): 2041–2064.
- Miravete, Eugenio J., Katja Seim, and Jeff Thurk.** 2018. “Market Power and the Laffer Curve.” *Econometrica*, 86(5): 1651–1687.
- Miravete, Eugenio J., Katja Seim, and Jeff Thurk.** 2020. “One Markup to Rule Them All: Taxation by Liquor Pricing Regulation.” *American Economic Journal: Microeconomics*, 12(1): 1–41.
- Nevo, Aviv.** 2001. “Measuring Market Power in the Ready-to-Eat Cereal Industry.” *Econometrica*, 69(2): 307–342.
- Office for National Statistics, Department for Environment, Food and Rural Affairs.** 2007-2012. “Living Costs and Food Survey, 2012.” *3rd Edition. UK Data Service. SN: 7472, DOI: http://doi.org/10.5255/UKDA-SN-7472-3*.
- Polyakova, Maria, and Stephen Ryan.** 2019. “Subsidy Targeting with Market Power.” *NBER Working Paper 26367*.

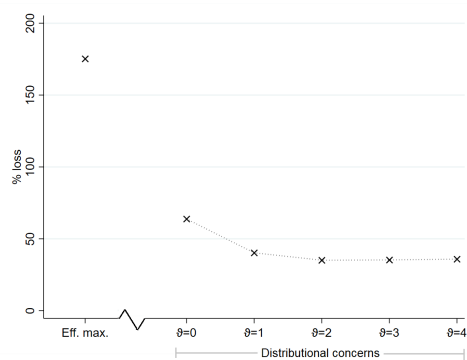
- Preonas, Louis.** 2019. "Market Power in Coal Shipping and Implications for U.S. Climate Policy."
- Roberto, Christina A., Hannah G. Lawman, Michael T. LeVasseur, Nandita Mitra, Ana Peterhans, Bradley Herring, and Sara N. Bleich.** 2019. "Association of a Beverage Tax on Sugar-Sweetened and Artificially Sweetened Beverages With Changes in Beverage Prices and Sales at Chain Retailers in a Large Urban Setting." *JAMA*, 321(18): 1799.
- Rojas, Christian, and Emily Yucai Wang.** 2017. "Do Taxes for Soda and Sugary Drinks Work? Scanner Data Evidence from Berkeley and Washington." Available at SSRN: 3041989.
- Saez, Emmanuel, and Gabriel Zucman.** 2016. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." *Quarterly Journal of Economics*, 131(2): 519–578.
- Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz.** 2012. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature*, 50(1): 3–50.
- Scientific Advisory Committee on Nutrition.** 2015. "SACN Carbohydrates and Health Report." <https://www.gov.uk/government/publications/sacn-carbohydrates-and-health-report>.
- Seiler, Stephan, Anna Tuchman, and Song Yao.** 2021. "The Impact of Soda Taxes: Pass-through, Tax Avoidance, and Nutritional Effects." *Journal of Marketing Research*, 58(1): 24–49.
- Smith, Matthew, Owen Zidar, and Eric Zwick.** 2020. "Top Wealth in America: New Estimates and Implications for Taxing the Rich." 117.
- Spiegler, Ran.** 2010. *Bounded Rationality and Industrial Organization*. Oxford University Press.
- Syverson, Chad.** 2019. "Macroeconomics and Market Power: Context, Implications, and Open Questions." *Journal of Economic Perspectives*, 33(3): 23–43.
- Tebaldi, Pietro.** 2017. "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA." *University of Chicago Working Paper*.
- Villas-Boas, Sofia Berto.** 2007. "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data." *Review of Economic Studies*, 74(2): 625–652.
- Wang, Emily Yucai.** 2015. "The Impact of Soda Taxes on Consumer Welfare: Implications of Storability and Taste Heterogeneity." *RAND Journal of Economics*, 46(2): 409–441.

- Wang, Y. Claire, Pamela Coxson, Yu-Ming Shen, Lee Goldman, and Kirsten Bibbins-Domingo.** 2012. "A Penny-Per-Ounce Tax On Sugar-Sweetened Beverages Would Cut Health And Cost Burdens Of Diabetes." *Health Affairs*, 31(1): 199–207.
- Weyl, E Glen, and Michal Fabinger.** 2013. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, 121(3): 528–583.
- World Health Organization.** 2015. *Sugar Intake for Adults and Children*. Geneva.



(a) Optimal tax rate

(b) Welfare gain



(c) Loss from ignoring market power

FIGURE 4. IMPACT OF DIFFERENT PREFERENCES FOR EQUITY ON OPTIMAL SIN TAX POLICY

Note: The horizontal axis plots social preferences. Eff. max. corresponds to an efficiency-maximizing government. Under distributional concerns the social marginal welfare weights on overseas individuals are 0 and the weights on domestic individuals are $(\hat{y}_i)^{-\theta}$. Panel (a) shows how the optimal tax rate (p/10oz) varies, panel (b) shows how the gain in welfare (per annum) relative to no drinks tax varies and panel (c) shows how the % welfare loss from policymaking that ignores market power varies.