Griliches Lectures

Lecture 1

Rachel Griffith

Singapore, July 2022

Zvi Griliches

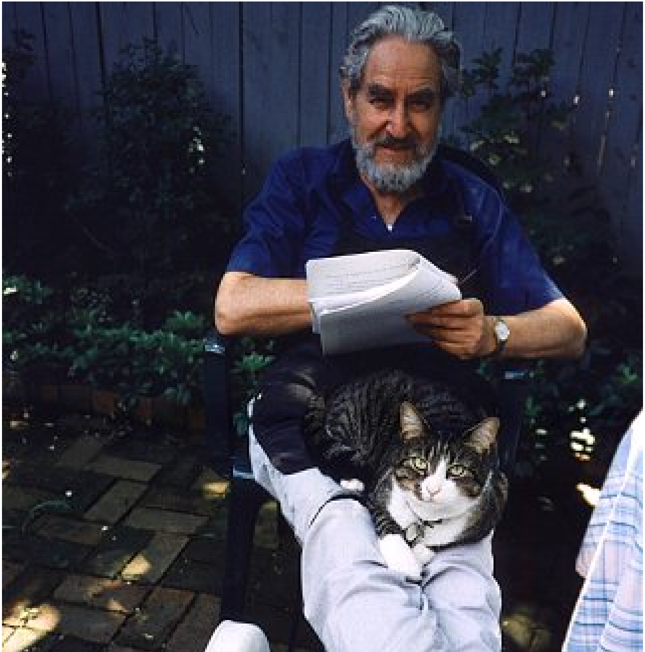

Motivation

Policy has the potential to improve welfare when consumption generates social costs; applies to "sin goods" - tobacco, alcoholic drinks, sugary drinks and other unhealthy foods (such as fast food)

- this motivates corrective taxes and regulations to availability, advertising, reformulation
- social costs include
  - externalities: second hand smoking, alcohol related crime, healthcare costs due to rising obesity
  - internalities: poor performance at school, future poor health, worse social and economic outcomes
- role of policy is to discourage socially costly consumption

Motivation

How effective are policies

- what the direct and indirect effects of policies?
- do they discourage socially costly consumption?
- what are the welfare implications? who gains and who loses?
- could we design better policies, that discourages socially cost consumption at lower cost?

Empirical approaches to learn about these questions ex ante

Plan of the lectures

Lecture 1

- what are the effects of sin taxes?
- empirical approaches to estimate suitably flexible demand models

Lecture 2

- what are the effects of restrictions to advertising?
- empirical approaches to estimate suitably flexible demand models
- evaluating welfare with possible behavioural effects

Lecture 3

- how do taxes and advertising interact?
- empirical approaches to learn about supply side dynamics

These lectures draw heavily on work with

Pierre Dubois (Toulouse)	Martin O'Connell (Wisconsin)	Kate Smith (LSE)
		

How well targeted are sin taxes?

How effective are taxes on alcoholic drinks and sugar sweetened beverages at discourage socially costly drinking?

- What is the distribution of social costs across the population of consumers?
- How do people respond to price change? And how does this vary across the population?
- How do firms respond to the tax?

Could we design a better targeted policy?

Griffith, O'Connell and Smith (2019) "Tax design in the alcohol market" Journal of Public Economics, 172, April 2019, 20-35

Griffith, O'Connell and Smith (2022) "Price floors and externality correction" Economic Journal, 123: 646, 2273-2289

Dubois, Griffith and O'Connell (2020) "How well targeted are soda taxes?" American Economic Review, 110 (11), November 202

O'Connell and Smith (2023) "Optimal sin taxation and market power" American Economic Journal: Applied Economics

Corrective taxes

- Pigou (1920): if the marginal externality is constant and equal across consumers then a tax can fully correct for the externality.
- However, marginal externalities often vary
  - may be nonlinear in quantity consumed
  - conditional on quantity, may vary across people
- Diamond (1973) considers the case of heterogeneous marginal externalities and a homogeneous good:
  - a linear tax can no longer achieve the first best
  - optimal tax rate equals weighted average marginal externality
- what about the case with heterogeneous marginal externality and differentiated goods?

Application to alcohol taxes

- existing alcohol taxes are poorly targeted
  - there are many differentiated products and the products that consumers who generate a lot of externalities like most face the lowest tax rate
- consumers' product demand curves are correlated with their marginal externality from ethanol consumption, which means:
  - higher tax rates on alcohol products disproportionately preferred by high marginal externality consumers allows policy to specifically target high externality generating consumption
- the potential for this will depend on the shape of demand - the price responsiveness of different types of consumers, and how correlates with marginal externality

Application to alcohol taxes

Griffith, O'Connell and Smith (2019):

- how important is this idea in practice in the UK alcohol market?
- what is the shape of demand?
- how does it differ across people, particularly how does it correlate with their marginal externality

How do we estimate the shape of demand

see Nevo (2000) on Readings, https://www.rachelgriffith.org/teaching

"Standard" early demand model (a la Deaton, 1986):

$\mathbf{q} = D(\mathbf{p}, \mathbf{z}, \epsilon)$

$\mathbf{q}$: vector of quantities
$\mathbf{p}$: vector of prices
$\mathbf{z}$: exogenous demand shifters, ϵ: random shocks
vectors are $J \times 1$ where J is number of products

Not feasible for markets considered here

- dimensionality problem – J is large, so too many parameters
- hard to deal with zeros
- not easy to flexibly accommodate heterogeneity in preferences

Models in Characteristics Space

Solution is to treat products as bundles of characteristics over which consumers have preference (Gorman (1956, 1980), Lancester (1966), Berry, Levinsohn and Pakes (1995), Nevo (2000))

- some products are better substitutes for each other than others
- rather than group products in ad hoc ways the characteristics define their substitutability
- reduces the dimensionality from number of products to number of characteristics
- key challenge is how to deal with unobservable characteristics

Usually implemented in a discrete choice random utility model (McFadden (1974, 1978, 1980, 1984))

- but does not have to be (e.g. Dubois, Griffith and Nevo (2014))

Models of differentiated demand

Consumer i in market t solves the choice problem

$$\max_{\{0,1,\dots,J\}} U(\mathbf{x}_{jt}, y_i - p_{jt}, \epsilon_{ijt}; \theta_i)$$

- $j = \{1, \dots, J\}$: products
- $\mathbf{x}_{jt}$: product characteristics
- y_i: consumer income
- p_{jt}: product price
- ϵ_{ijt}: idiosyncratic shock to utility
- θ_i: consumer's preferences

specify functional form of utility $U()$ and consumer preferences, θ_i

this will impose important restrictions on shape of demand

Importance of consumer heterogeneity

without heteogeneity (observed and unobserved) in consumer preferences, price elasticities are:

$$\eta_{jkt} = \frac{\partial s_{jt}}{\partial p_{kt}} \frac{p_{kt}}{s_{jt}} = \begin{cases} -\alpha p_{jt}(1 - s_{jt}) & \text{if } j = k \\ \alpha p_{kt} s_{kt} & \text{otherwise} \end{cases}$$

- own price elasticities: market shares are typically small, so $\alpha(1 - s_{jt})$ is nearly constant and therefore the own-price elasticities are proportional to price
- cross-price elasticities: cross price elasticity wrt a change in the price of product k is that same for all products such that $j \neq k$

The effects of allowing heterogeneity

see BLP (1995) on Readings, https://www.rachelgriffith.org/teaching

- Heterogeneity in tastes for product attributes drives correlation
  - for example, if "luxury" is an attribute of a car, then a consumer who likes one luxury car is more likely than the average consumer to like another luxury car
- Price elasticities

$$\eta_{jkt} = \frac{\partial s_{jt}}{\partial p_{kt}} \frac{p_{kt}}{s_{jt}} = \begin{cases} -\frac{p_{jt}}{s_{jt}} \int \alpha_i s_{ijt}(1 - s_{ijt}) dP_D(D) dP_v(v) & \text{if } j = k \\ \frac{p_{kt}}{s_{jt}} \int \alpha_i s_{ijt} s_{ikt} dP_D(D) dP_v(v) & \text{otherwise} \end{cases}$$

  - own-price elasticity no longer driven by functional form; e.g. will depend on price sensitivity of consumers who are attracted to that product
  - cross-price elasticities no driven by a priori segmentation, and also quite flexible

Empirical model of alcohol demand

Why a discrete choice demand model:

- switching between disaggregate alcohol products
  - common in the literature on alcohol taxes to aggregate products into a relatively small number of categories (e.g. beer, wine etc.)
  - but this masks the considerable variation in price and alcoholic strength within category
- correlation between product demands and total ethanol demand
- avoids the curse of dimensionality
- rationalises zero purchases
- well suited for incorporating rich preference heterogeneity

Discrete choice demand model

On each choice occasion consumer selects from 69 alcohol options or an outside option (no purchase)

- embed decision to purchase alcohol or not, what product to choose and what quantity

Observed and unobserved preference heterogeneity:

- household specific preferences can vary over price, alcohol strength, ethanol content and market segment (random coefficients)
- vary by {light, moderate, heavy} drinker $\times$ {low, mid, high} income

Why allow preferences to vary with drinking habits?

Alcohol consumption is concentrated - literature suggests that marginal externalities are as well

(a) UK

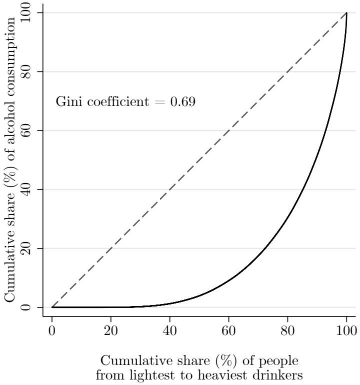

(b) US

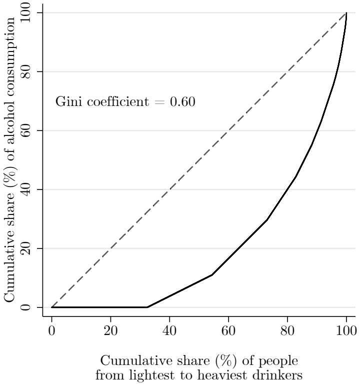

Source: UK Health Survey for England 2011; US National Health and Nutrition Survey 2007–2012.

Variation in the marginal externality/internality

Evidence suggests that externalities are convex in alcohol consumption

- the more you drink the greater the external cost associated with one more drink
  - threshold effect with some diseases: at low levels of alcohol consumption the risk is not elevated, but this risk increases sharply above a certain point
  - higher levels of alcohol consumption create an exponential risk of accidents: odds of injury from 8 pints almost 18 times greater than the odds of injury from 1 pint
- It could also be that some people generate more externality (for any given level of drink) than others

Demand specification

Utility household i gets from option (j, s) in period t:

$$u_{ijst} = \nu(p_{jst}, z_{js}, \mathbf{x}_{jst}; \theta_i) + \epsilon_{ijst}$$

where

$$\nu(p_{jst}, z_{js}, \mathbf{x}_{jst}; \theta_i) = \alpha_i p_{jst} + \beta_i w_j + \sum_{m=1}^{4} 1[j \in \mathcal{M}_m] \cdot (\gamma_{i,1m} z_{js} + \gamma_{i,2m} z_{js}^2) + \xi_{ijt}.$$

Unobserved product characteristic:

$$\xi_{ijt} = \eta_{ij} + \zeta_{k_j t}$$

preference parameters $\alpha_i, \beta_i, \gamma_i$ are normally distributed across consumers, with the possibility that households have correlated tastes for price, alcoholic strength and ethanol content

Data

Household/individual level scanner data increasingly available

- prices, quantities and characteristics at the household and product level
- longitudinal (follow the same households over time)
- facilitates modeling consumer preference heterogeneity
- Berry and Haile (2016), consumer level data allow more flexibility (relative to market level data) for identification of differentiated products demand models
- repeated observations for the same individual over time, allows us to capture the degree of correlation in a decision maker's choices, as the prices and product characteristics they face change, provides information about the strength of their individual preferences

UK alcohol market - data

Use household scanner data from the Kantar Worldpanel

- over 11,500 households observed buying alcohol 2010-2011
  - estimate demand on 2011 data
  - use 2010 as pre-sample information to identify persistently heavy drinkers

  There are 7000+ alcohol UPCs (barcodes) and 3000 brands
  - aggregate to 32 "products"
    - available in different sizes, 69 product-sizes
  - focus on the margins of substitution that are most relevant to our application
  - these capture heterogeneity in the shape of demand for sets of UPCs that are impacted similarly by alcohol tax changes

Social planner's problem

Social planner chooses:

- a vector of tax rates (levied per unit of ethanol), $\boldsymbol{\tau}$

to maximise:

- the sum of consumer surplus and tax revenue minus the externality:

$$\max_{\boldsymbol{\tau}} W(\boldsymbol{\tau}) = \underbrace{\sum_i \left[y_i + \frac{v_i(\boldsymbol{\tau})}{\alpha_i} \right]}_{\text{consumer surplus}} + \underbrace{R(\boldsymbol{\tau})}_{\text{tax revenue}} - \underbrace{\Phi(\boldsymbol{\tau})}_{\text{external costs}}$$

Alternative tax policies

Consumer specific taxes: set the rate equal to the marginal externality:

$$\tau_i^* = \phi_i'(Z_i(\tau_i^*))$$

This achieves the first best (Pigou, 1920)

Single ethanol tax rate: the optimal rate is:

$$\tau^* = \underbrace{\bar{\phi}'}_{\substack{\text{Average}\\ \text{marginal}\\ \text{externality}}} + \underbrace{\frac{\text{cov}(\phi_i', |Z_i|))}{|\bar{Z}'|}}_{\substack{\text{covariance of the marginal externality}\\ \text{and slope of ethanol demands}}}$$

where $Z_i' = \sum_j z_j \frac{\partial q_{ij}}{\partial \tau}$ (Diamond, 1973)

Alternative tax policies

Product level tax rates: the optimal tax rates will solve the first order conditions for each τ_j:

$$\sum_i \sum_k (\tau_k - \phi_i') \frac{\partial Z_{ik}}{\partial \tau_l} = 0$$

When there is a correlated heterogeneity in marginal externalities and demands, there are welfare gains from varying tax rates across products

All else equal, the tax rate on a set of products is:

1. higher if it is highly demanded by high marginal externality consumers
2. higher the stronger is the correlation between the marginal externality and the own-price elasticity
3. lower the stronger the correlation between the marginal externality and cross slopes of demand

Externality function

Assume the externality function, ϕ, takes the form:

$$\phi(Z_{it}) = \phi_0(\exp(\phi_1 Z_{it}) - 1)$$

where $Z_{it} = \sum_j z_j q_{ijt}$ denotes derived ethanol demand.

Calibrate:

- ϕ_0 to match UK government estimate of aggregate external costs
- ϕ_1 determines how concentrated the costs are

Show how results vary with aggregate external costs and degree of convexity...

... and with heterogeneity in (ϕ_0, ϕ_1) across socioeconomic groups.

Current UK alcohol taxes

Figure 1. Duty per unit of alcohol

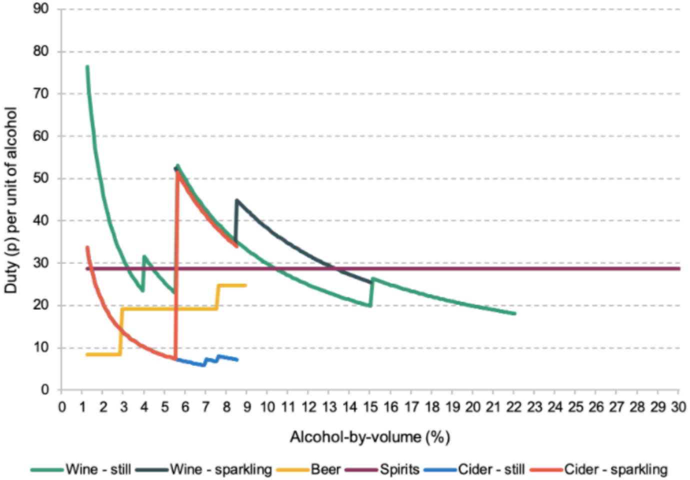

Source: Kate Smith (2021), https://ifs.org.uk/publications/15761

Optimal single rate

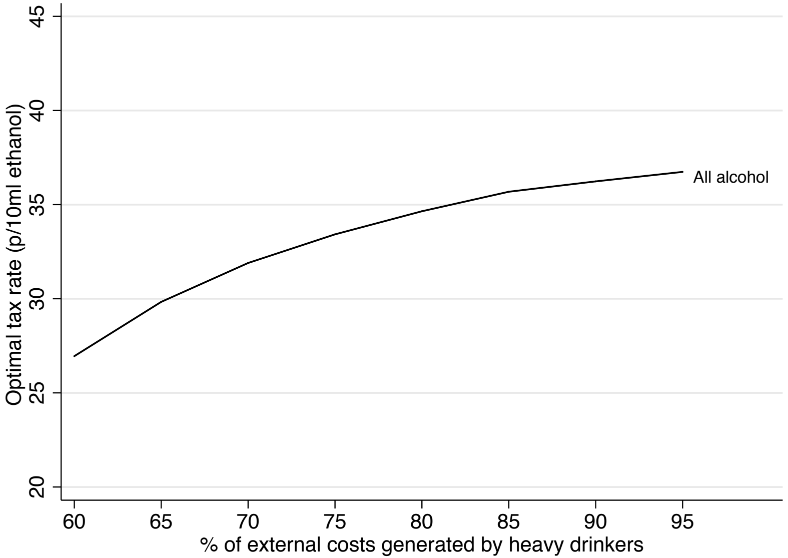

Optimal multi rate

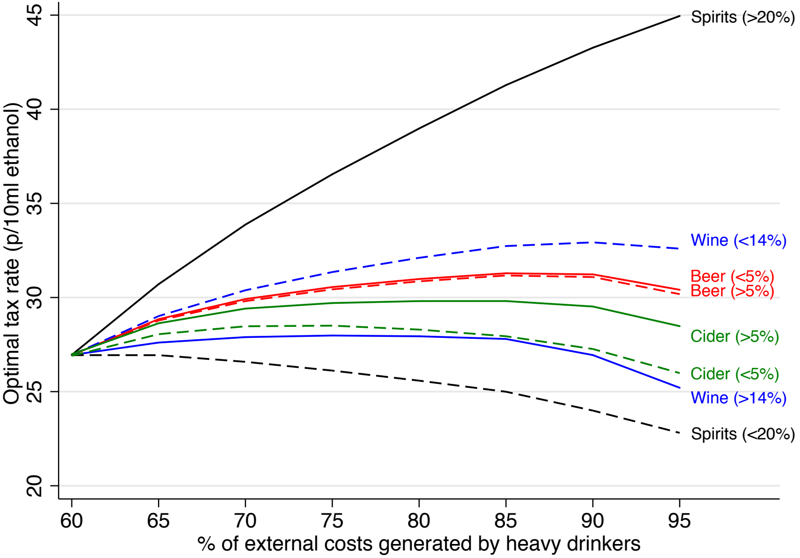

Notes: The figures show the optimal tax rates under various calibrations of the convexity of the externality function, shown on

Welfare

the sum of consumer surplus and tax revenue minus external costs

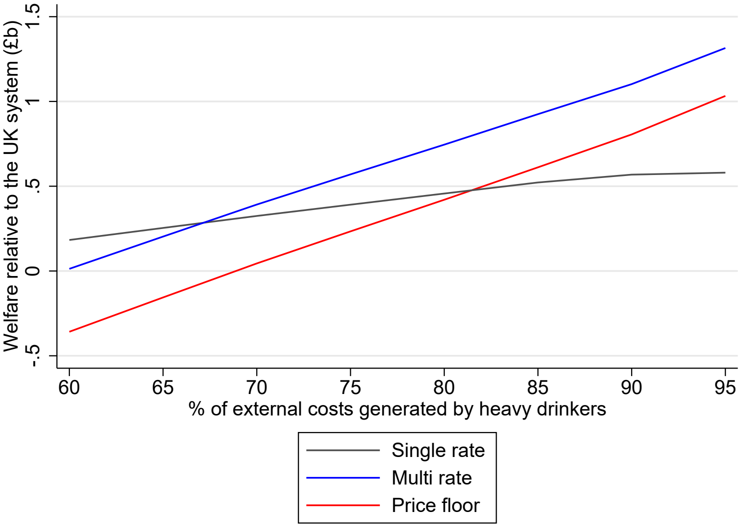

Summary and conclusions

Consider corrective tax design to correct consumption externalities in markets in which:

- marginal externalities vary across consumers
- there are many differentiated products
- there is potential correlation between marginal externalities and shape of demand, allowing the possibility to better target externality generating consumption

These ideas have empirical relevance in the UK alcohol market:

- an optimal system that varies rates across alcohol types would improve welfare relative to a single rate, if the externality function is convex

Soda taxes

Dubois, P., R. Griffith and M. O'Connell (2020) "How well targeted are soda taxes?" American Economic Review, 110 (11), November 2020

- exploit the long time dimension of panel scanner data to estimate consumer specific preference parameters (in contrast to treating them as random draws from a flexible distribution)
- exploits data collected on individuals and therefore avoids making implicit assumptions about how households' choice relates to individual welfare

Soda taxes - motivation

Sugar consumption well above recommended level

- excess consumption creates externalities (higher health and social care costs)
- and internalities (worse long term health, social and economic outcomes)

Of particularly concern for children in low income households

Policies aim to reduce externality/internality generating consumption

- where they raise the price of non-externality generating consumption this incurs a welfare loss
- where they raise the price of internality generating consumption and don't change behaviour this incurs a welfare loss

Externality/internality function

Externality/internality function unknown, we assume:

- young
  - likely higher, and consequences higher, in the young
- high total annual dietary sugar
  - possible convexity (e.g. probability of developing type II diabetes rises non-linearly in sugar consumption)
- income
  - the stress and cognitive load of being in poverty means more likely to make unwise decisions and underweight the future

Why soda taxes?

Sugar intensity of calories (%)	% of added sugar from soft drinks	% of added sugar from confectionery
< 5	8.2	7.6
5-10	13.8	9.6
10-15	17.4	12.4
15-20	26.0	13.2
> 20	35.1	16.3

Source: UK National Diet and Nutrition Survey 2008–2011.

Dubois, Griffith and O'Connell (2020)

- choice behavior of individuals (rather than households)
- on-the-go accounts for 45% of sugar from soft drinks, represents higher fraction for young people
- non-parametric specification of preferences
  - exploit panel nature of data, which has sufficiently long time series
  - allows us to relate preferences to any individual attributes or outcome in a flexible way
  - avoids potentially restrictive distributional assumptions
- data on individuals and the household they live in
  - Data follow 5,500 individuals through time
  - Use household level data to estimate at-home demand
  - Firm pricing also depends on the at-home segment

The non-alcoholic drinks market

- Data from Kantar On-The-Go Survey over 2009–2014
  - 5,554 individuals followed through time
  - records all food and drinks purchases made in shops, vending machines and kiosks
  - includes details of transaction price, outlet and precise product
  - over 25 drinks purchased, for over half of individuals over 75 purchases
- Data from Kantar Worldpanel
  - covers grocery purchases made and brought into home by 4,204 household including households to which individuals belong

The non-alcoholic drinks market

Model behavior in the non-alcoholic drinks market, which includes

- Soft drinks (sugary and diet)
  - e.g. 330ml can of Coca Cola
- Alternative sugary drinks
  - e.g. fruit juice, flavoured milk
- Bottled water

Market demand from on-the-go and at-home demands

- estimate demand in both segments of the market
- aggregate across segments to obtain market demand

Demand model

- Choice occasion τ refers to consumer purchasing a drink in retailer r_τ at time t_τ
- Payoff consumer i from product j on choice occasion τ:

$$U_{ij\tau} = \underbrace{\alpha_i p_{jr_\tau t_\tau} + \beta_i s_j + \gamma_i w_j}_{v_{ijr_\tau t_\tau}} + \underbrace{\delta^z_{d(i)} z_j + \delta^h_{d(i)} h_{c(i)t_\tau} + \xi_{d(i)b(j)t_\tau} + \zeta_{d(i)b(j)r_\tau}}_{\eta_{ijr_\tau t_\tau}} + \epsilon_{ij\tau}$$

- Preferences over price (p_{jrt}), sugar (s_j), and soft drink (w_j) are consumer specific
- Other preferences are demographic group specific: other attributes ($\mathbf{z}_j$), brand-time (ξ_{dbt}), brand-retailer (ζ_{dbr}) effects, temperature h_{ct} in county c at t

Preference heterogeneity

- Estimate $(\alpha_i, \beta_i, \gamma_i)$ by consumer i exploiting large $\mathcal{T}$
  - Allow infinite preference in $(\alpha_i, \beta_i, \gamma_i)$ space
    - Consumer always(never) purchase inside products (soft drinks) has $\gamma_i = +\infty(-\infty)$
    - Consumers always(never) select sugary options has $\beta_i = +\infty(-\infty)$.
  - Yields individual choice sets that include/exclude soft drinks, sugary varieties, diet varieties; identified by individual purchases (large T)
- Other preferences by demographic group
  - on-the-go: four groups based on gender and age
  - at-home: five groups based on household composition

Demand model

- With $\epsilon_{ij\tau}$ iid type I extreme value, choice probability i purchases good j on occasion τ:

$$P_{i\tau}(j) = \frac{1_{\{j \in \Omega_i\}} \exp(v_{ijr_\tau t_\tau} + \eta_{ijr_\tau t_\tau})}{\sum_{k \in \Omega_i \cap \Omega_{r_\tau}} \exp(v_{ikr_\tau t_\tau} + \eta_{ikr_\tau t_\tau})}$$

  - Ω_{r_τ} denotes the consumer's choice set on that occasion
- Estimate using Maximum likelihood on many choice occasions by individual (large T)
  - Avoid assumptions of standard random coefficient approach
  - Robustness to incidental parameters problems (Hahn and Newey, 2004) using split sample jackknife (Dhaene and Jochmans, 2015)

Identification

- Two sources of price variation
  - Cross-retailer variation in *relative* prices and availability of products conditional on brand-time and retailer-drink type effects
    - assume retailer choice not driven by shocks to demand for a *specific* drinks product
  - Variation in prices within brand across containers and sizes.
    - assume no aggregate shocks within brand for different container types
- Tilting of price schedules driven by cost variation not proportional to pack size and differential pass-through of cost shocks
- Variation in sugar content between sugary and diet varieties
  - assume brand effects common across sugary and diet varieties, and taste for sugary varieties additively separable

Demand Estimates - On-the-go

Moments of distribution of consumer specific preferences			
Variable		Estimate	Standard error
Price (α_i)	Mean	-3.0737	0.0287
	Standard deviation	2.6825	0.0210
	Skewness	-0.9247	0.0462
	Kurtosis	4.3175	0.1117
Soft drinks (γ_i)	Mean	1.4297	0.0421
	Standard deviation	1.6065	0.0153
	Skewness	0.5001	0.0415
	Kurtosis	3.6833	0.1307
Sugar (β_i)	Mean	0.4244	0.0104
	Standard deviation	1.8058	0.0141
	Skewness	-0.4838	0.0407
	Kurtosis	3.5801	0.1026
Price-Soft drinks	Covariance	-1.4058	0.0463
Price-Sugar	Covariance	0.7413	0.0439
Soft drinks-Sugar	Covariance	-0.6585	0.0311
Demographic specific carton-size effects ($\delta^z_{d(i)}$)		Yes	
Demographic specific weather effects ($\delta^h_{d(i)}$)		Yes	
Time-demographic-brand effects ($\xi_{d(i)b(j)t}$)		Yes	
Retailer-demographic-brand effects ($\zeta_{d(i)b(j)r}$)		Yes	

Demand Estimates - On-the-go

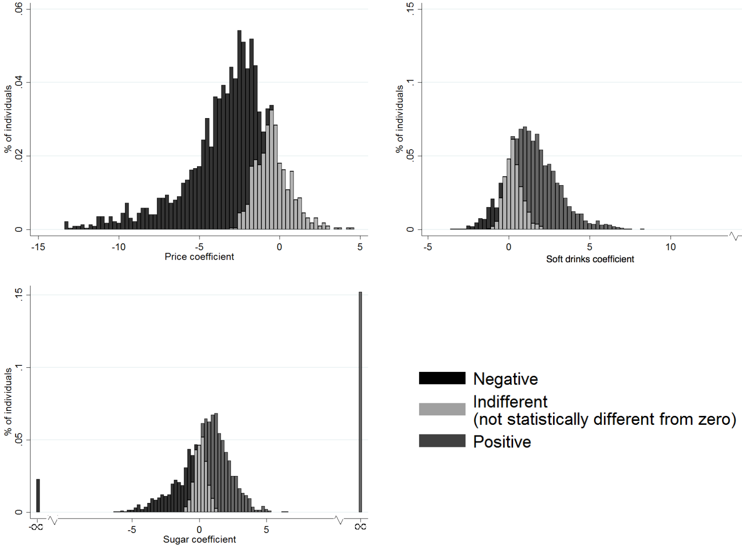

Demand Estimates - On-the-go

Variation in sugar preferences with age

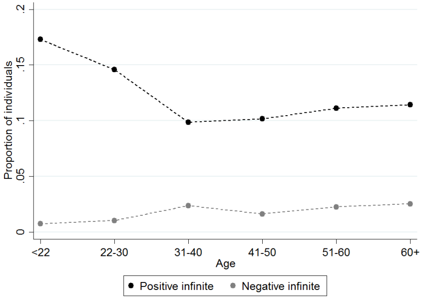

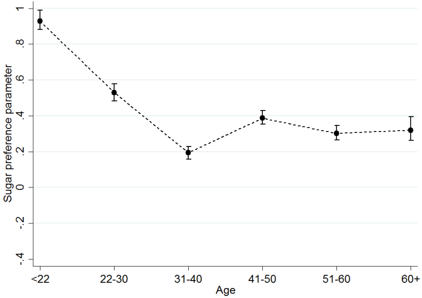

Demand Estimates - On-the-go

Variation in price preferences with age

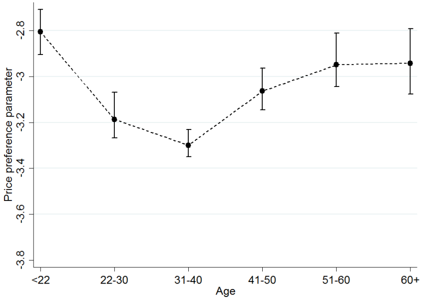

Price preference

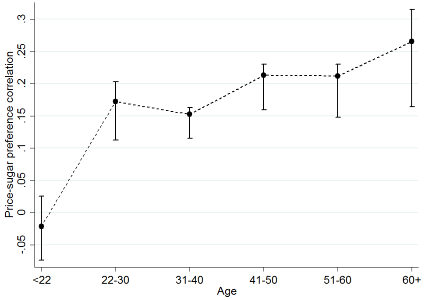

Sugar-price preference correlation

Supply model

- Demand for product j in market m depends on consumer demand on-the-go ($\mathcal{M}_m^{out}$) and at-home ($\mathcal{M}_m^{in}$) segments:

$$q_{jm}(\mathbf{p}_m) = \underbrace{M_m^{out} \sum_{(\tau,i)\in\mathcal{M}_m^{out}} P_{i\tau}(j)}_{\equiv q_{jm}^{out}(\mathbf{p}_m)} + \underbrace{M_m^{in} \sum_{(\tau,i)\in\mathcal{M}_m^{in}} P_{i\tau}(j)}_{\equiv q_{jm}^{in}(\mathbf{p}_m)}$$

- Firm f's (variable) profits in market m are given by:

$$\Pi_{fm} = \sum_{j\in F_f} (p_{jm} - c_{jm}) q_{jm}(\mathbf{p}_m)$$

and the firm's price first order conditions are:

$$q_{jm}(\mathbf{p}_m) + \sum_{k\in F_f} (p_{km} - c_{km}) \frac{\partial q_{km}(\mathbf{p}_m)}{\partial p_{jm}} = 0 \quad \forall j \in F_f.$$

allowing us to estimate marginal costs $\mathbf{c}_m$

Tax pass-through

- Tax creates wedge between pre ($\tilde{\mathbf{p}}$) and post ($\mathbf{p}$) tax prices,
- Volumetric tax, π, on sugary soft drinks implies:

$$p_{jm} = \begin{cases} \tilde{p}_{jm} + \pi l_j & \forall j \in \Omega_{ws} \\ \tilde{p}_{jm} & \forall j \in \Omega_{wd} \bigcup \Omega_{as} \bigcup \Omega_{an} \end{cases}$$

  where l_j is the volume of product j

- Counterfactual equilibrium, prices satisfy:

$$q_{jm}(\mathbf{p}_m) + \sum_{k \in F_f} (\tilde{p}_{km} - c_{km}) \frac{\partial q_{km}(\mathbf{p}_m)}{\partial p_{jm}} = 0 \quad \forall j \in F_f$$

- Takes into account overlap in products across segments

Aggregate effect of tax

	Tax (pence)	Δ price (pence)	Δ share (p.p.)
Sugary soft drinks	10.65	13.15	-6.34
			[-6.41, -6.24]
Diet soft drinks	0.00	-1.37	3.96
			[3.92, 4.01]
Sugary alternatives	0.00	0.00	1.09
			[1.06, 1.13]
Outside option	0.00	0.00	1.28
			[1.25, 1.30]

Tax of 25 pence per litre on sugary soft drinks (similar to California and UK)

How well targeted is the tax?

Reductions in sugar from drinks, by age and total sugar consumption

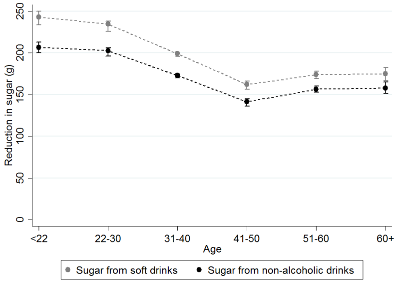

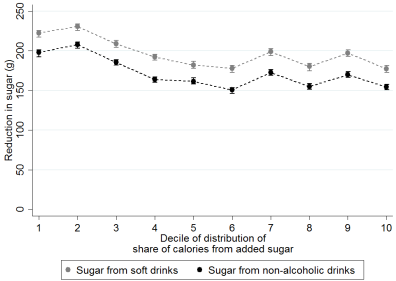

How well targeted is the tax?

Reductions in sugar from drinks, by age and total sugar consumption

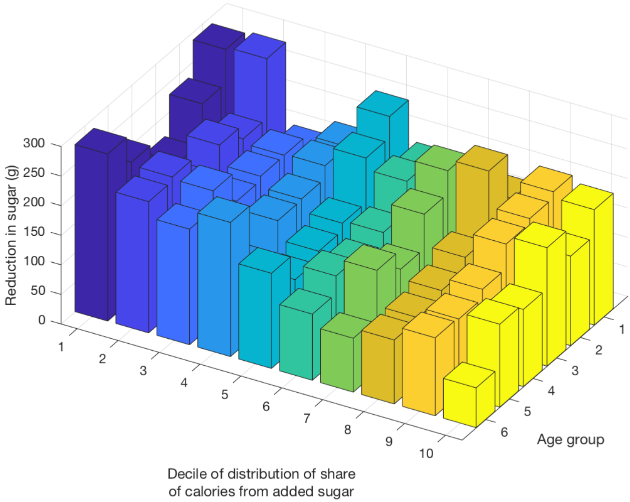

Reduction in sugar from drinks

Summary

- Larger reductions amongst the
  - young
    - have strong preferences for sugar, but are responsive to price
  - low income
  - young and low-income
  - not those with high overall sugar consumption
    - they having strong sugar preferences and are relatively price insensitive

Compensating variation

by age and total sugar consumption

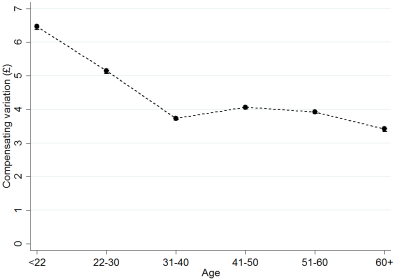

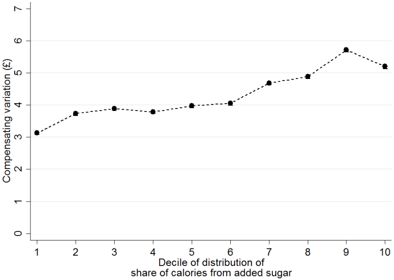

Compensating variation

by age and total sugar consumption

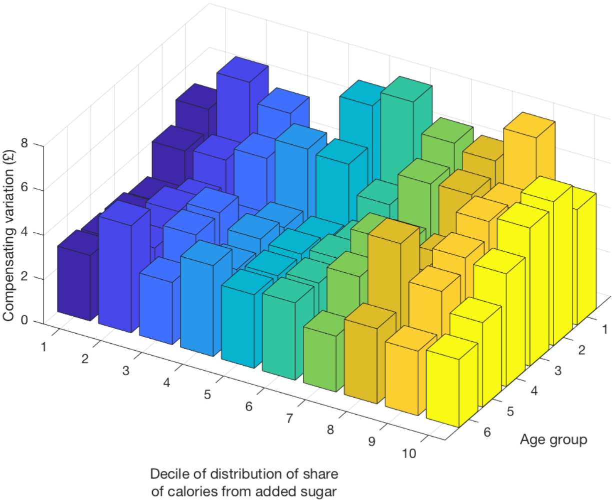

Compensating variation

Summary

- Compensating variation is higher for the young, low income and high overall sugar consumers
- Internalities mean this is only part of the total consumer welfare effect
- Average compensating variation for young soft drink purchasers is £6.47, and average reduction in sugar is 207g
- If internality associated with drinking the amount of sugar in a can of Coca Cola is above £1.10, then for the average young person the soda tax will be welfare improving
- If tax revenue, which is £3.56 per person, is redistributed lump-sum to soda purchasers then this threshold would be £0.49 per can of Coca Cola

Conclusion

Results point towards a (sugary) soda tax

1. Relatively effective at lowering sugar intake of younger consumers
   - because they buy more of the taxed products
   - and respond more strongly in *level* terms
2. Less well targeted at lowering sugar intake of those with systematically high levels of dietary sugar
   - despite being relatively heavy consumers of sugary soft drinks
   - strong sugar preferences and price insensitivity leads this group to respond less strongly
3. Results in somewhat larger "revealed" consumer welfare losses among lower incomes, but also achieves larger sugar reductions among them