Institute for Fiscal Studies

MANCHESTER 1824

Recent advances in Public Economics: using demand estimates to inform policy

Rachel Griffith

RES Masterclass, October 2024

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 of goods, advertising, and reformulation
- social costs include
  - externalities (costs imposed on others): e.g. second hand smoking, alcohol related crime, healthcare costs due to rising obesity
  - internalities (costs imposed on yourself in future): e.g. lower future health, worse social and economic outcomes
- role of policy is to discourage socially costly consumption

Motivation

How effective are policies

- what are the direct and indirect effects of policies?
- does policy lead to reduction in socially costly consumption?
- how does supply respond?
- what are the welfare implications? who gains and who loses?
- could we design better policies, that discourages socially cost consumption at lower cost?

I will discuss empirical approaches to learn about these questions ex ante

Plan of the lectures

Lecture 1

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

Lecture 2

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

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 that led to greater reductions in social costs and/or imposed less costs?

Draws on these papers

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

see Readings, https://www.rachelgriffith.org/teaching

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

- People are heterogeneous in the social costs they generate from drinking alcohol
- Alcohol consumption is concentrated

(a) UK

(b) US

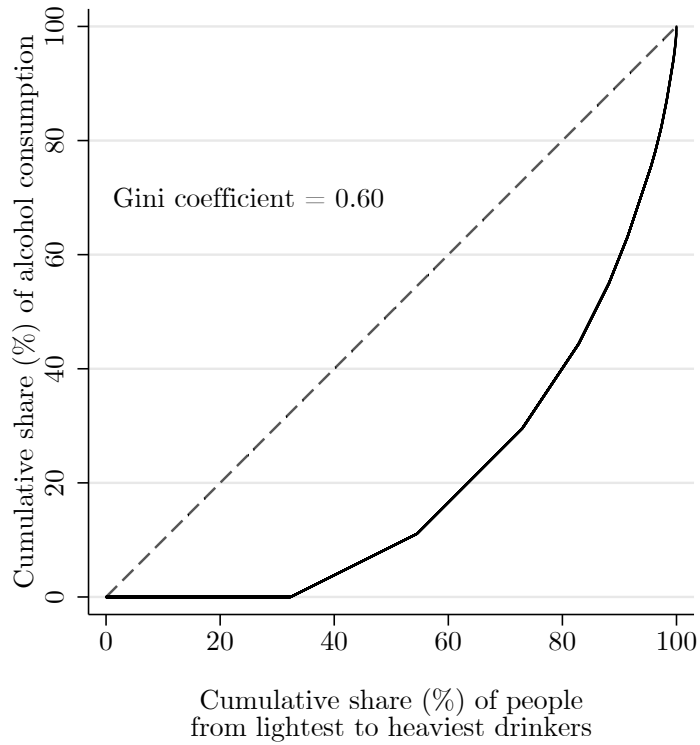

Source: UK Health Survey for England 2011; US National Health and Nutrition Survey 2007–2012.

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

Application to alcohol taxes

- products are heterogeneous
  - there are many differentiated products, that vary by alcohol strength, and by other characteristics
- existing alcohol taxes are poorly targeted, products that consumers who generate a lot of externalities like face the lowest tax rate

Figure 1. Duty per unit of alcohol

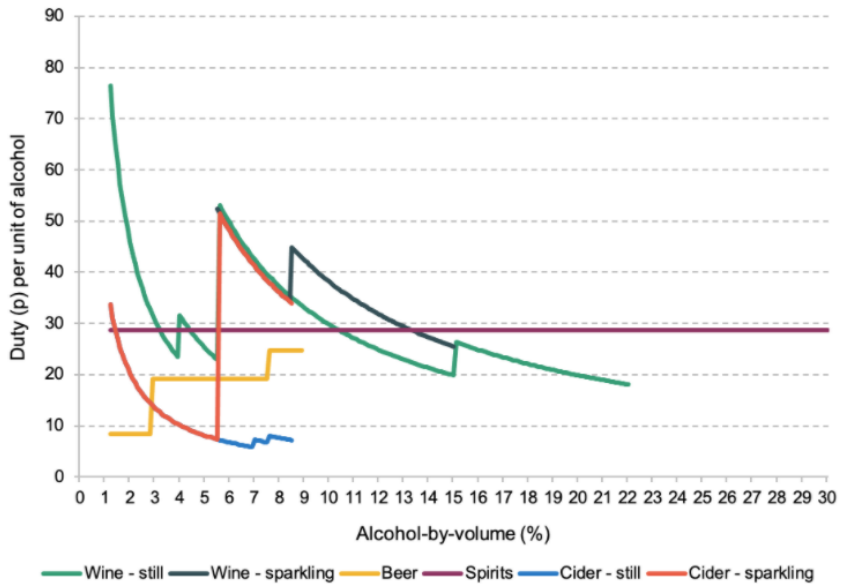

Source: Kate Smith (2021), https://ifs.org.uk/publications/15761

Application to alcohol taxes

- how important is this idea in practice in the UK alcohol market?
- what is the shape of demand?
- if consumers' product demand curves are correlated with their marginal externality from ethanol consumption, this 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
  - specifically, what is the price responsiveness of different types of consumers, and how does it correlates with the marginal externality they generate

How do we estimate the shape of demand

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

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

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

$\mathbf{q}$: vector of quantities
$\mathbf{p}$: vector of prices
$\mathbf{z}$: exogenous demand shifters, $\boldsymbol{\epsilon}$: 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

Typically implemented in a discrete choice random utility model

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)
- preferences vary by $\{$light, moderate, heavy$\}$ drinker $\times$ $\{$low, mid, high$\}$ income

Demand specification

Utility household i gets from product j, of size s in period t:

$$u_{ijst} = \overbrace{\alpha_i p_{jst}}^{price} + \underbrace{\beta_i w_j}_{strength} + \overbrace{(\gamma_{i0} + \gamma_{d(i)m(j)1}) z_{js} + \gamma_{d(i)m(j)2} z_{js}^2}^{ethanol} + \underbrace{\xi_{d(i)jt} + \eta_{im(j)}}_{unobserved} + \epsilon_{ijst}$$

- 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
- $d(i)$ is heterogeneity across demographic group ($\{$light, moderate, heavy$\}$ drinker $\times$ $\{$low, mid, high$\}$ income)
- $m(j)$ is heterogeneity across product segment (premium low strength beer, ..., red wine, ..., premium gin, ...)

Data

Household/individual level scanner data

- prices, quantities, characteristics at 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:

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

where $Z_i(\tau_i^*)$ is ethanol demand.

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 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:

$$\phi(Z_{it}) = \phi_0 Z_{it} + \phi_1 Z_{it}^2$$

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

- (ϕ_0, ϕ_1) determine the aggregate external cost and degree of convexity of the function
- we calibrate to match UK government estimate of aggregate external costs

We show how results vary with aggregate external costs and degree of convexity

Optimal single rate

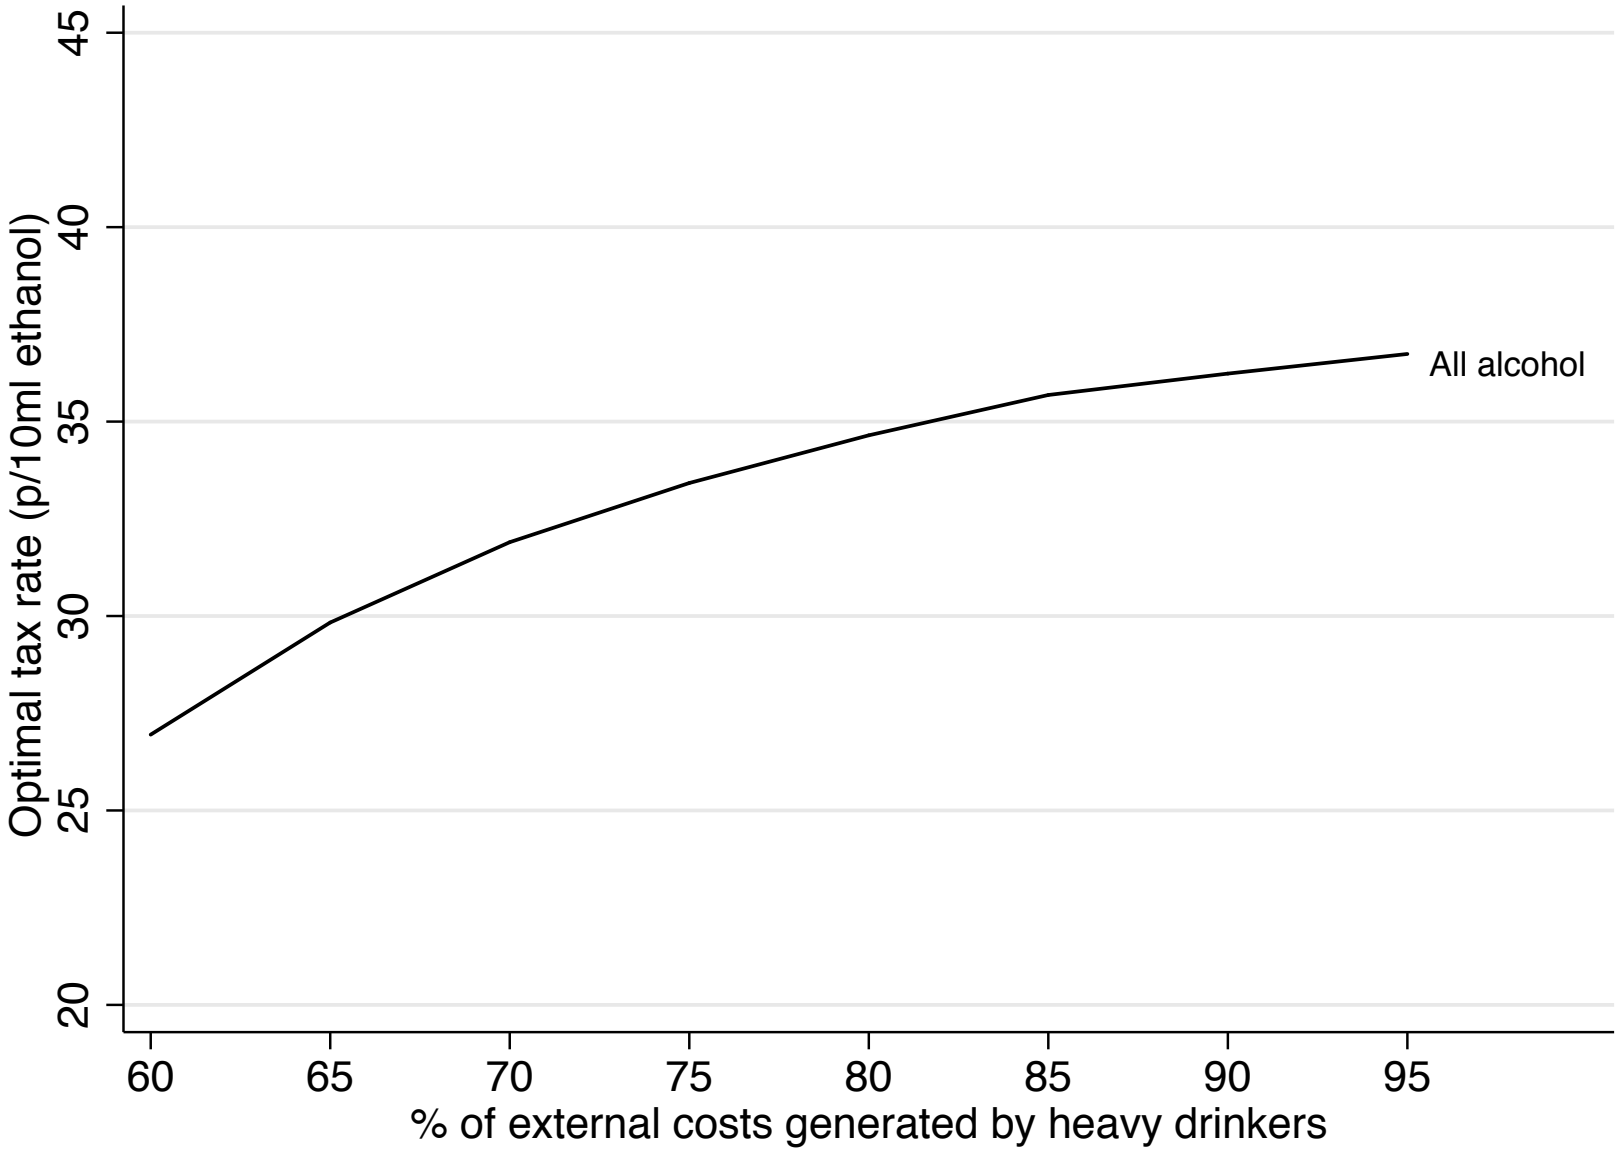

Optimal multi rate

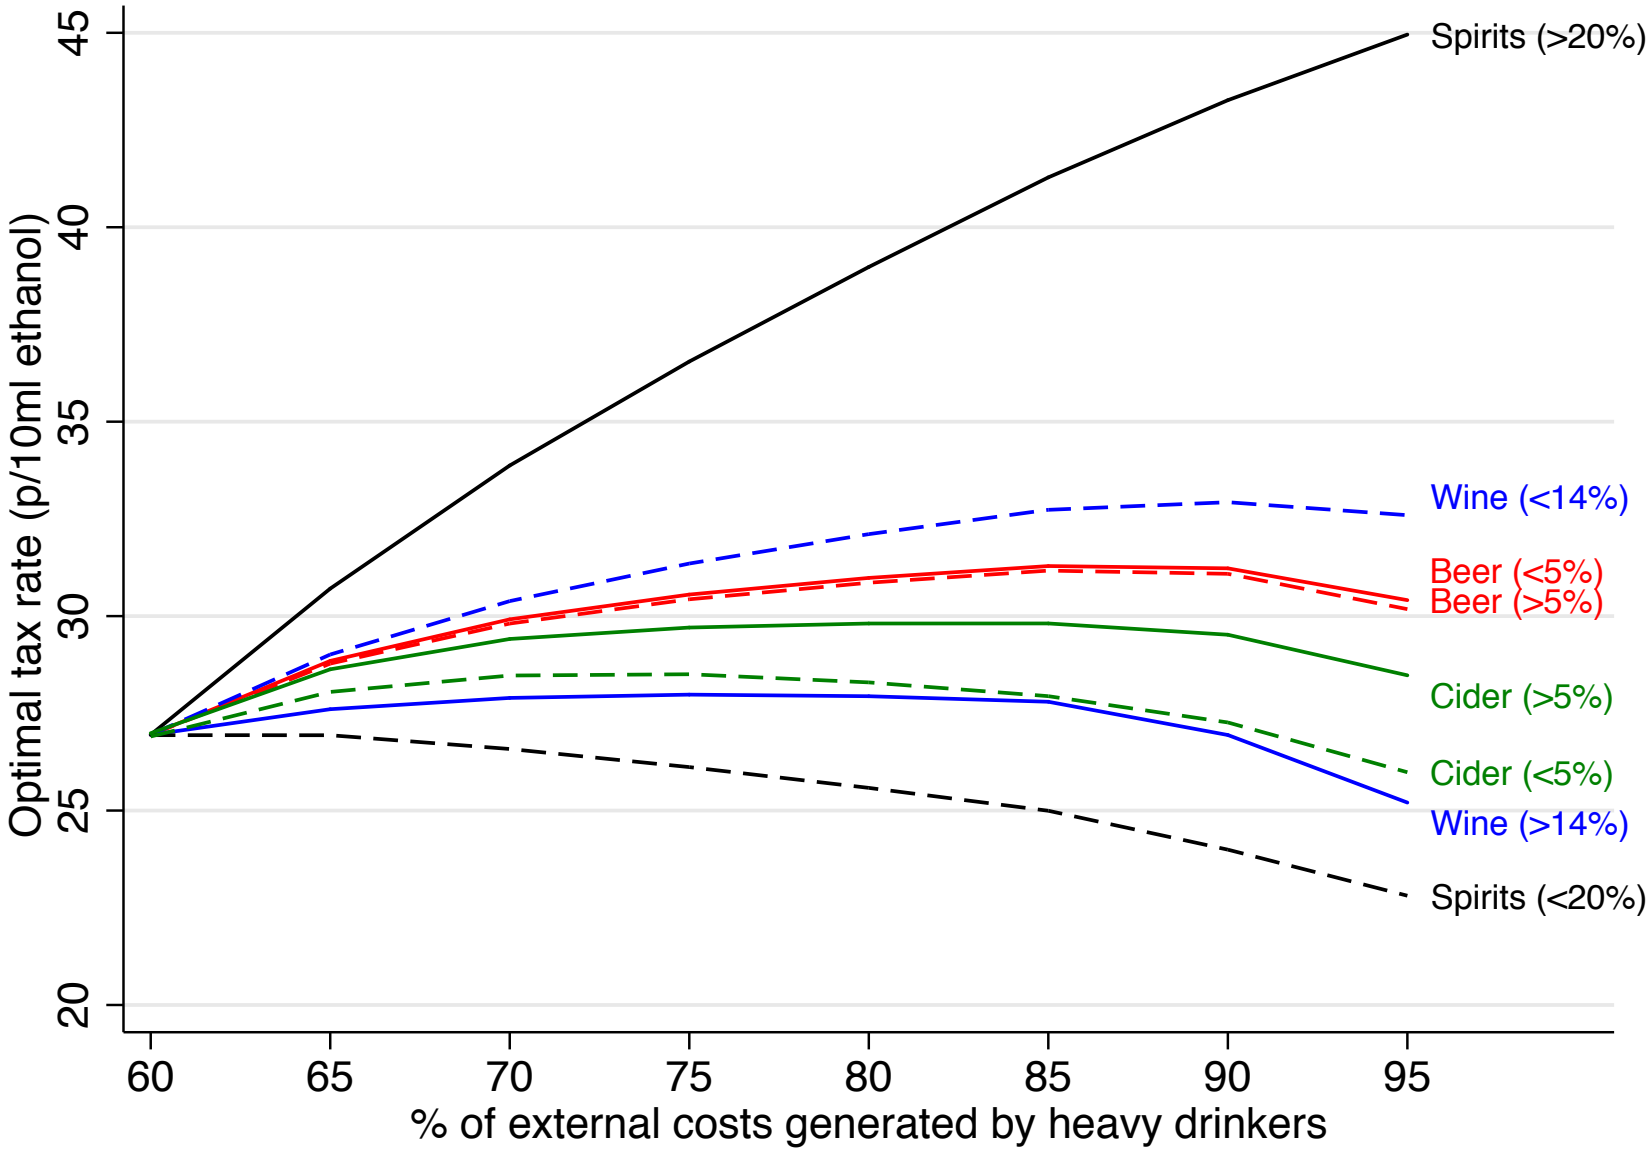

Notes: The figures show the optimal tax rates under various calibrations of the convexity of the externality function, shown on

Welfare

the sum of consumer surplus and tax revenue minus external costs

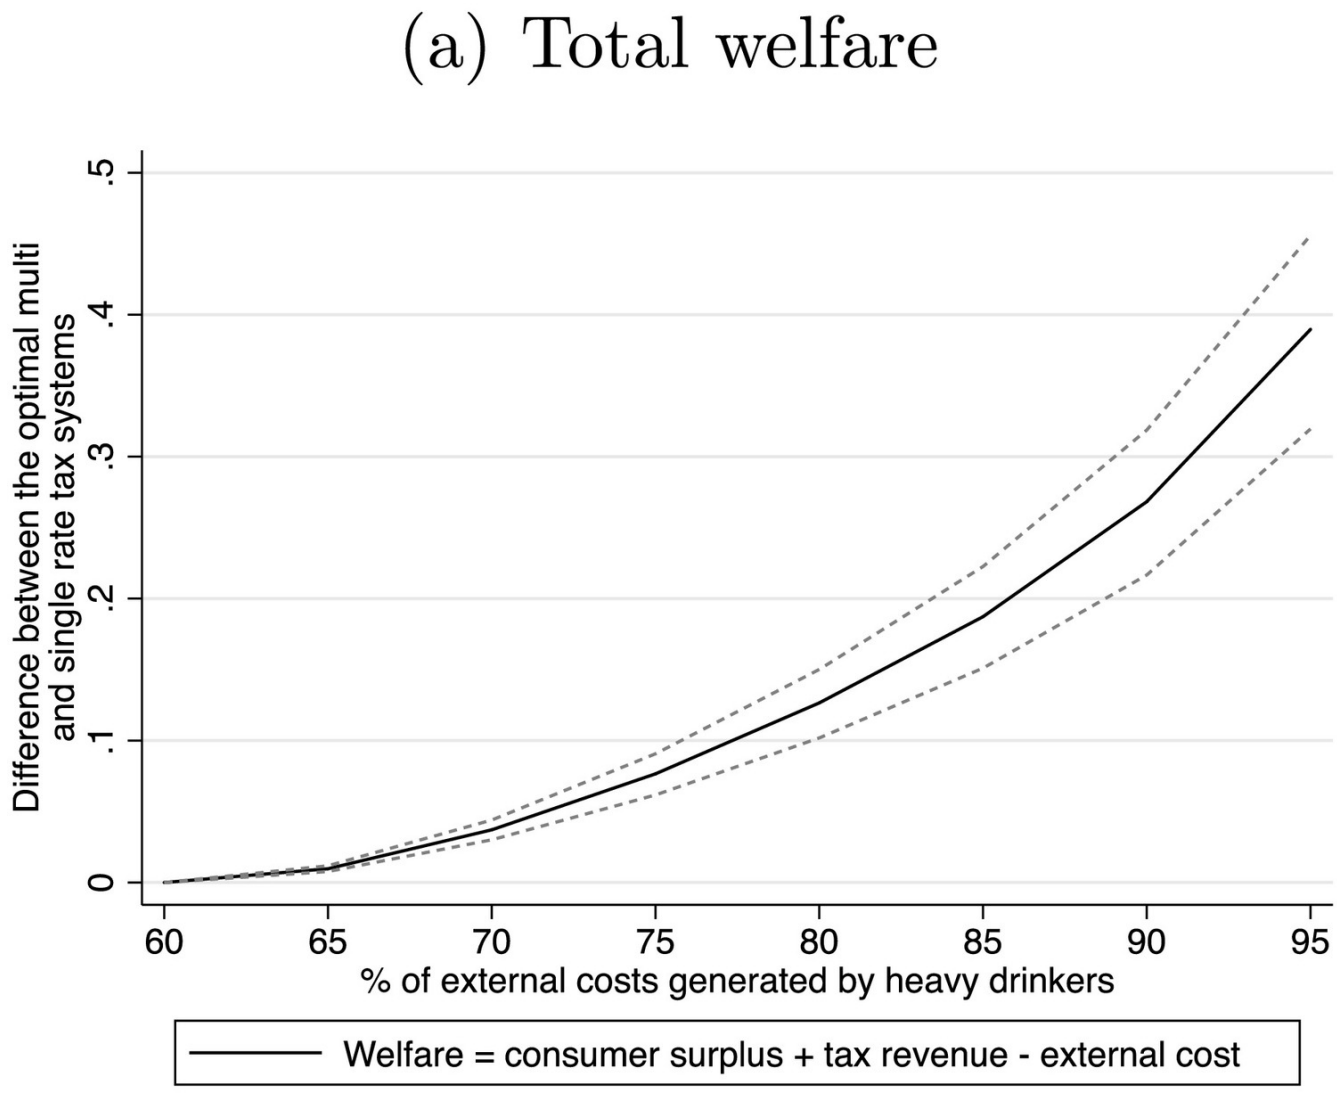

Current UK system of alcohol taxes

(a) Current UK system

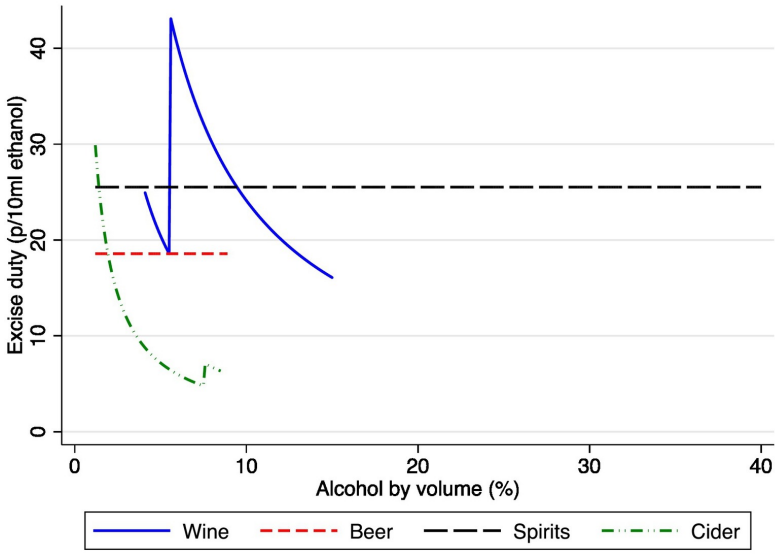

(b) Optimal UK system

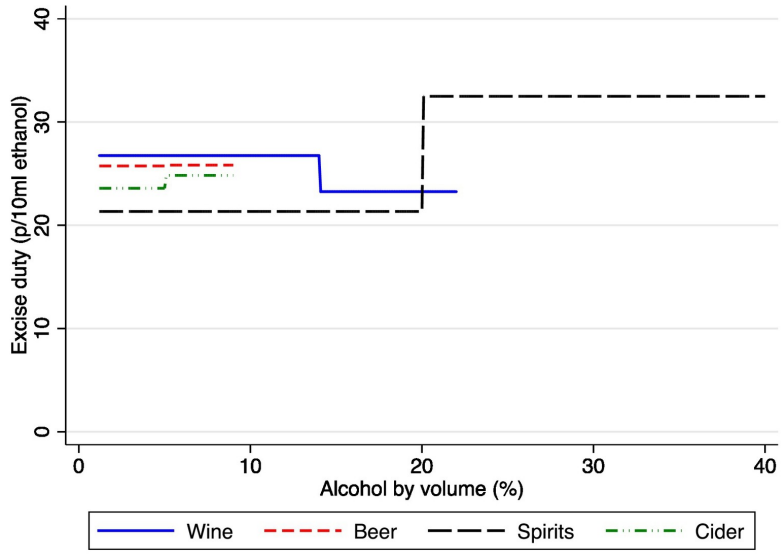

Summary and conclusions

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

Policy relevance

The Head of Excise Branch at HMT said

> *"I would like to highlight in particular your work on the optimal design of alcohol taxes. The IFS report ... directly contributed to the Chancellor's decision both to launch a review of the alcohol duty system at the 2020 Budget and to focus the review on economic rationality.*
> *For individual academic papers to be highlighted in this way is unusual and a testament to their quality and high level of policy relevance."*

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, but we believe:

- 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

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:

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

- Preferences over price (p_{jrt}), sugar (s_j), soft drink (w_j) consumer specific
- unobserved characteristics: brand ($\mathbf{z}_j$), brand-time (ξ_{dbt}), brand-retailer (ζ_{dbr}) effects, temperature h_{ct} in regional (county) c and time t level

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 Estimates - On-the-go

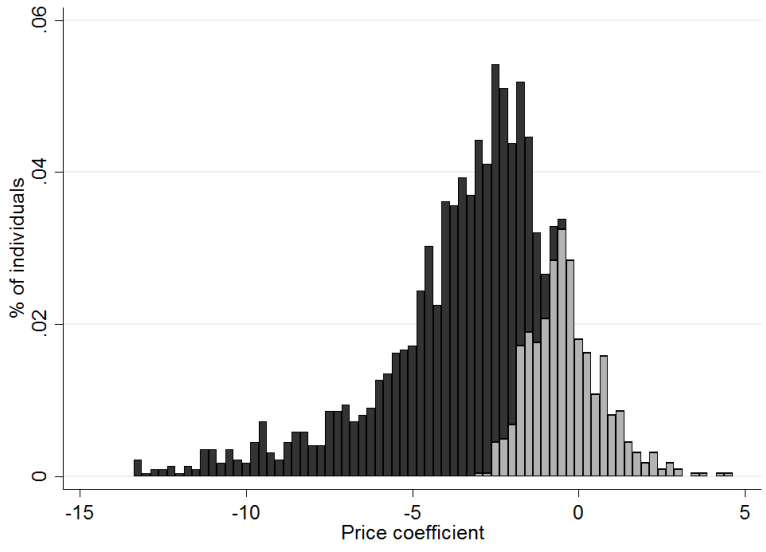

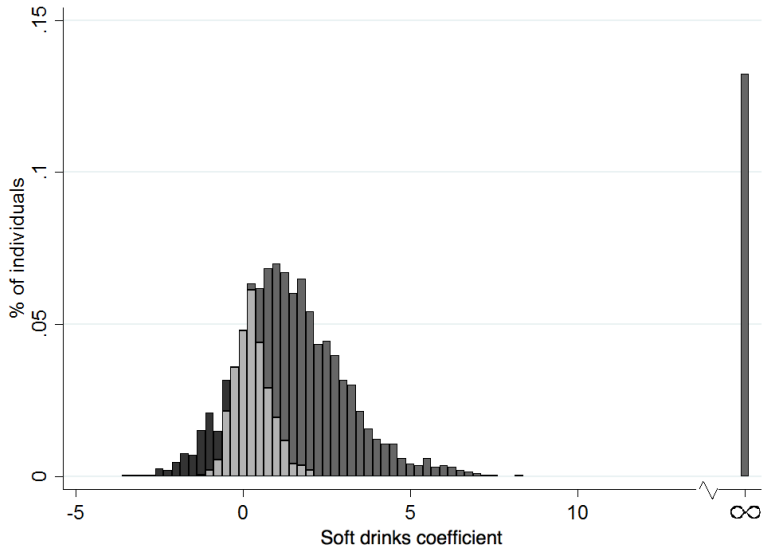

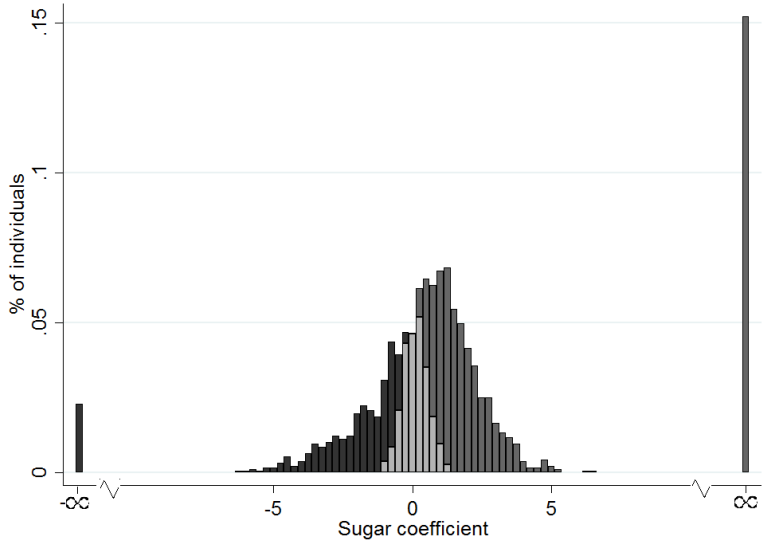

Negative

Indifferent
(not statistically different from zero)

Positive

Demand Estimates - On-the-go

- Variation in sugar preferences with age
  - young people like sugar more

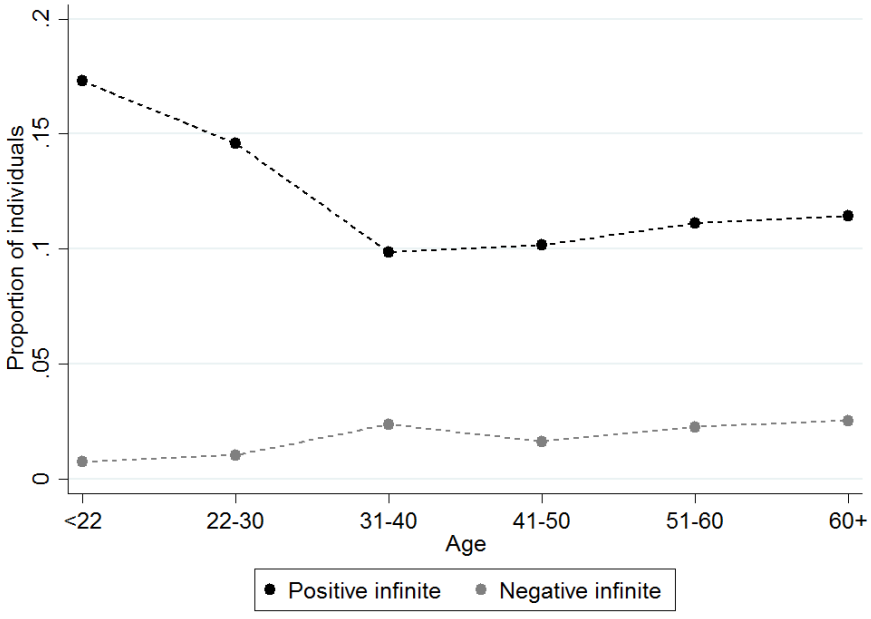

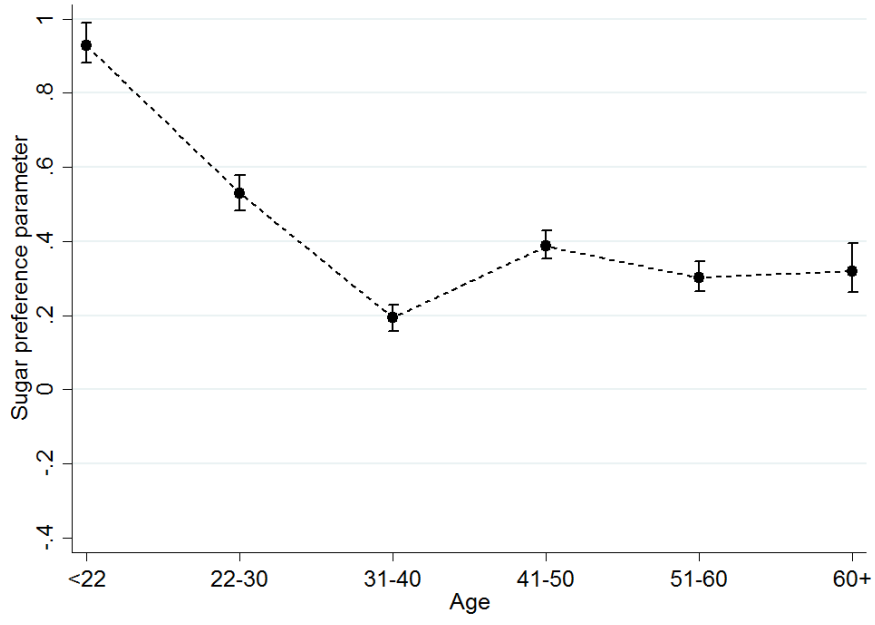

Demand Estimates - On-the-go

- Variation in price preferences with age
  - young people are (a bit) less sensitive to price
  - for young people, the ones who like sugar more are also more price sensitive, for older people the opposite is true

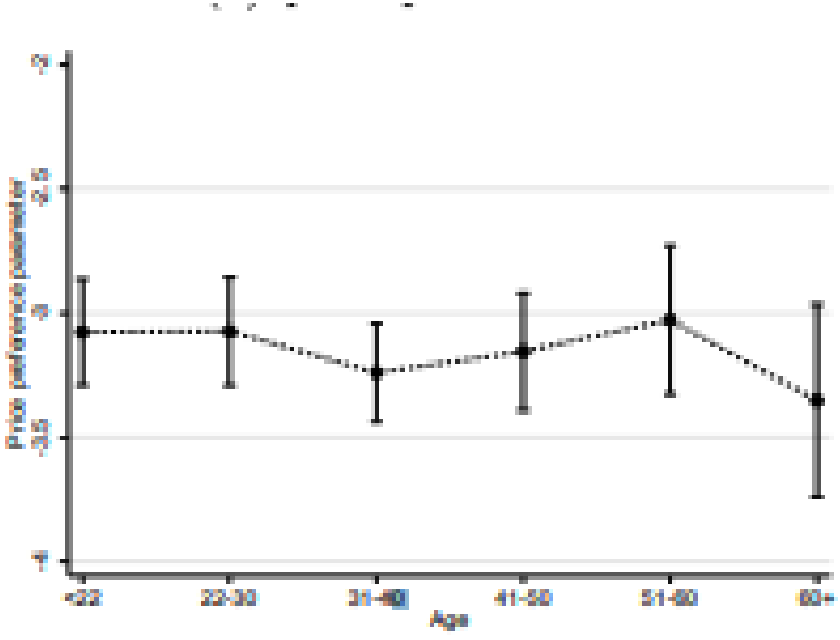

Price preference

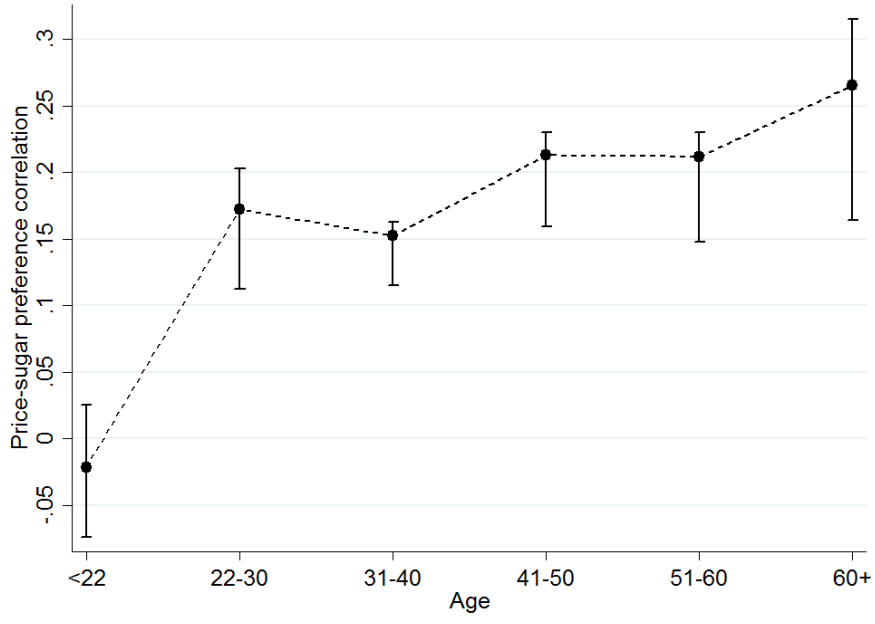

Sugar-price preference correlation

What effect does the tax have?

- We specify a supply model and estimate tax-pass through, consider a tax of 25 pence per litre on sugary soft drinks

	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]

How well targeted is the tax?

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

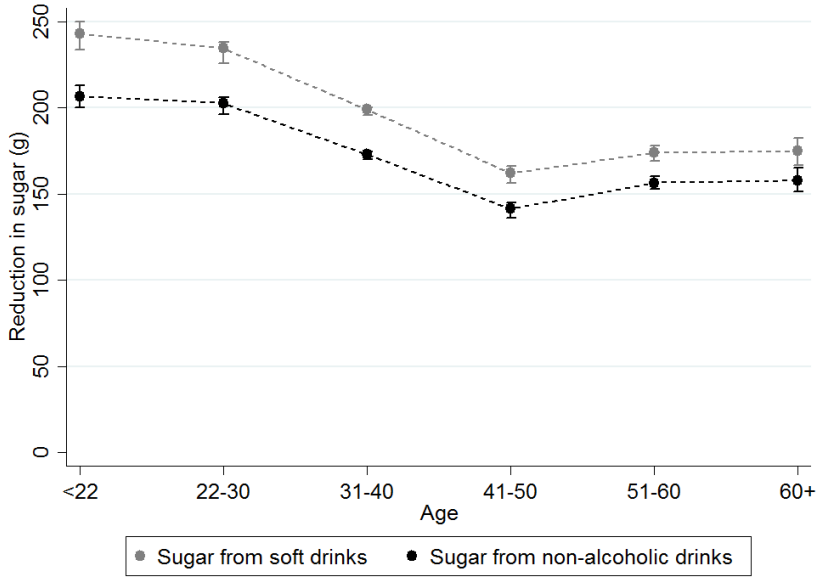

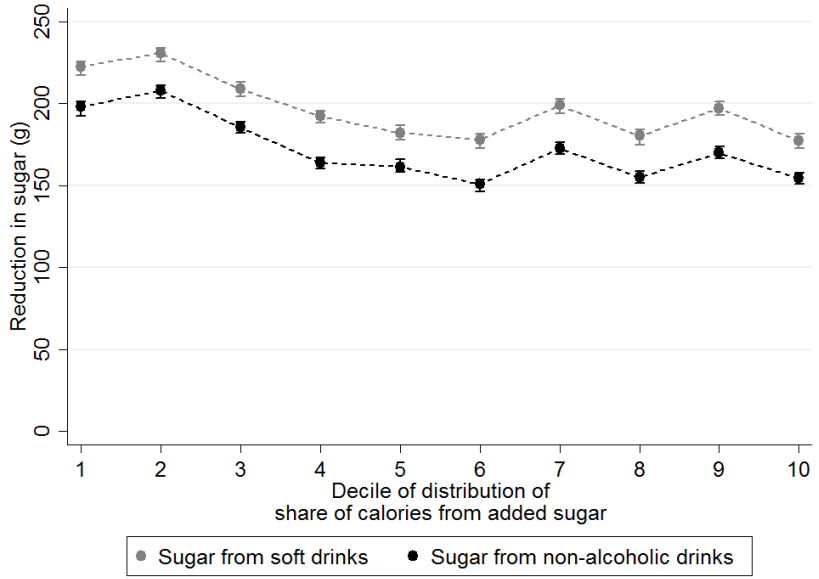

Reduction in sugar from drinks

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

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

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

Key points

- In order to empirically learn about the impact of a tax on reducing social costs it is essential that the demand model we estimate allows for sufficient flexibility in the shape of demand to allow for correlation between consumers' responses and the social costs they generate
- Exactly what flexibility is most important will depend on the application - both on how the market works, and what particular questions you are seeking to answer