Prices versus Preferences:

Taste Change and Revealed Preference

Abi Adams* Richard Blundell† Martin Browning*‡ Ian Crawford†§

August 2019¶

Abstract

A systematic approach for incorporating taste variation into a revealed preference framework for heterogeneous consumers is developed. This enables the recovery of the minimal variation in individual tastes required to rationalise observed choices. It is used to examine the extent to which changes in tobacco consumption are driven by price changes or by taste changes, and whether the significance of these two channels varies by socioeconomic factors. A censored quantile approach is developed that allows for unobserved heterogeneity and for individuals who choose not to consume tobacco. Statistically significant educational differences in the marginal willingness to pay for tobacco are recovered with more highly educated individuals experiencing a greater shift in their tastes.

*(Corresponding Author) Department of Economics, University of Oxford, Oxford OX1 3UQ, UK; abi.adams[at]economics.ox.ac.uk

†Department of Economics, University College London, Gower Street, London WC1E 6BT, UK; +44(0)20 7679 5863; r.blundell[at]ucl.ac.uk

‡Department of Economics, University of Oxford, Oxford OX1 3UQ, UK; martin.browning[at]economics.ox.ac.uk

§Department of Economics, University of Oxford, Oxford OX1 3UQ, UK; ian.crawford[at]economics.ox.ac.uk

¶We would like to thank Laurens Cherchye, Dennis Kristensen, Arthur Lewbel and participants at the CAM workshop, cemmap workshop and the AEA meetings session on consumer demand for helpful comments. The authors also gratefully acknowledge financial support from the Economic and Social Research Council Centre for the Microeconomic Analysis of Public Policy at IFS and the European Research Council grant MicroConLab. All errors are ours.

1 Introduction

Structural empirical research on consumer behaviour is typically based upon the idea of choice-revealed preference: consumers choose what they prefer from the options available to them and thereby reveal their preferences through repeated observations of their choices from different budget sets. Simple techniques can then be used to recover a consumer's preferences from data on their choices using methods developed by Samuelson (1948), Houthakker (1950), Afriat (1967) and Varian (1982). However, classical revealed preference methods can only be applied when preferences are *stable*. If preferences change during the period of observation then these methods cannot solve the inverse problem.

In this paper we develop a systematic approach to examine taste change within a revealed preference (RP) framework. We create a methodology that recovers the minimal variation in individual tastes that is required to rationalise observed choices. The patterns revealed by this exercise may suggest fruitful extensions to the deterministic taste model. Patterns of taste changes that appear to vary systematically over time may be diagnostic of, for example, unmeasured quality change, the spread of unobserved complimentary technologies, or information acquisition. Patterns of taste change that vary systematically over consumer characteristics may be evidence of preference heterogeneity. Patterns of taste variation which appear simply random may be diagnostic of mistakes or measurement errors.

We represent taste heterogeneity as a linear perturbation to an individual base utility function, much as in McFadden and Fosgerau (2012) and Brown and Matzkin (1998). Under this specification, taste change can be interpreted as the shift in the marginal utility of a good for each individual. We show that this specification is not at all restrictive. We derive inequalities that are an extension of Afriat (1967) such that when they hold there exists a well-behaved base utility function and a series of taste shifters that perfectly rationalise observed behaviour. We then show, under mild assumptions on the characteristics of available choice data, that we can always find a pattern of taste shifters on a single good that are sufficient to rationalise any finite time series of prices and quantities.

We apply to analyse the preferences for a good where there is strong *prima facie* evidence that tastes have changed: tobacco products. In particular, we ask how much of the fall in tobacco consumption is due to a rise in the relative price of tobacco and how much needs to be attributed to taste changes? We also consider how tastes evolve across different socio-economic strata, asking the question: Does education matter?

For the 1980s and 1990s in the UK (and elsewhere), when changes in tastes for tobacco consumption most likely occurred, there is no consumer panel data available. Instead we implement our approach on repeated cross-sections of household consumer expenditure survey data. We show that under a monotonicity restriction relating to unobserved heterogeneity, the demand for each individual consumer is identified by the conditional quantile applied to the repeated cross-section data (Blundell and Matzkin (2015)). In so doing there is a tight connection between the theory (which applies to individuals) and the data, in a setting in which we have sufficient price variation to reject the time-invariant tastes model. We also allow for censoring to account for consumers who choose not to consume tobacco. We apply the RP inequality conditions on the censored conditional quantile demand functions in order to recover the taste change for tobacco, and we also extend the analysis to allow for the non-separability of tobacco consumption from alcohol consumption.

Our objective is both better to understand the pattern of taste changes and to inform policy on the balance between information/health campaigns and tax reform. Governments have a limited set of levers should they wish to influence household consumption patterns. These include quantity constraints, price changes through taxes and subsidies and information programs. The relative efficacy of the different modes is important for designing public policy. The approach in this paper allows us to consider the extent to which changes in tobacco consumption are due to price changes and how much is due to preference change.

This paper proceeds as follows. Sections 2 outlines our theoretical framework and derives the necessary and sufficient conditions under which observed behaviour and our model of taste change are consistent. It then goes on to develop a quadratic programming methodology that can be applied to recover the minimal amount of interpersonally comparable taste variation that is necessary to rationalise choice behaviour. Section 3 introduces the data used for our empirical investigation of tobacco consumption in the UK and discusses the construction of quantity sequences that we construct from the UK Family Expenditure Survey using censored quantile regression methods. Section 4 applies our method to rationalise the changes in tobacco consumption occurring in the U.K. since 1980. Section 5 relaxes our empirical framework to permit tobacco and alcohol to be non-separable in the utility function. Finally, Section 6 concludes our analysis and considers the implications of our findings for government anti-smoking policy moving forward.

2 Theoretical framework

Consider a consumer who selects a quantity vector $\mathbf{q} \in \mathbb{R}^K_+$ at time t to maximise their time-dependent utility function:

$$u(\mathbf{q}; \boldsymbol{\alpha}_t) = v(\mathbf{q}) + \boldsymbol{\alpha}_t'\mathbf{q} \tag{1}$$

subject to $\mathbf{p}'\mathbf{q} = x$, where $\mathbf{p} \in \mathbb{R}^K_{++}$ is an exogenous price vector and x is total expenditure. It is assumed that $u(\mathbf{q}, \boldsymbol{\alpha}_t)$ is locally nonsatiated and concave conditional on $\boldsymbol{\alpha}_t \in \mathbb{R}^K$, where $\boldsymbol{\alpha}_t$ is a vector of marginal utility perturbations that indexes the consumer's tastes at time t. This utility function therefore consists of two components: a set of "base" preferences given by $v(\mathbf{q})$, and a time-varying part given by $\boldsymbol{\alpha}_t'\mathbf{q}$.[1] Our use of marginal utility shifters to capture taste changes follows the random utility approach of McFadden and Fosgerau (2012) and Brown and Matzkin (1998).

In general, taste changes of this kind can also be understood as innovations in the willingness-to-pay for goods and can therefore be given an equivalent interpretation in terms of virtual or support prices. Take a dataset made up of a sequence of price-quantity observations for a consumer $\{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,\dots,T}$. We can always find a sequence of taste-shifters $\{\alpha_t\}_{t=1,\dots,T}$ that can rationalise the model. To see this, consider the first order conditions

$$\nabla u(\mathbf{q}_t) = \nabla v(\mathbf{q}_t) + \alpha_t \leq \lambda_t \mathbf{p}_t$$

This can be rewritten in terms of the base preferences and shadow prices as

$$\nabla v(\mathbf{q}_t) \leq \lambda_t \tilde{\mathbf{p}}_t$$

where we use the shadow prices

$$\tilde{\mathbf{p}}_t = \mathbf{p}_t - \frac{\alpha_t}{\lambda_t}$$

where α_t/λ_t represents the innovations in willingness-to-pay or *'taste-wedge'* for every good in every period.

[1] The effect of the taste change parameters $\boldsymbol{\alpha}_t$ on consumer demand is not invariant to transformations of $v(\mathbf{q})$ and so further analysis is conditioned upon a given cardinal representation of the "base" preferences.

The behaviour generated by the model

$$\max_{\mathbf{q}} v(\mathbf{q}) + \boldsymbol{\alpha}_t'\mathbf{q} \text{ subject to } \mathbf{p}_t'\mathbf{q} = x_t$$

is therefore identical to the behaviour generated by the model

$$\max_{\mathbf{q}} v(\mathbf{q}) \text{ subject to } \tilde{\mathbf{p}}_t'\mathbf{q} = \tilde{x}_t$$

where preferences are not subject to taste change, but where the prices and budget are replaced by their virtual counterparts.[2] Thus the question of whether there exist rationalising taste-shifters is equivalent to the question of whether we can always find shadow prices which can rationalise the observed quantity data. Varian (1988) shows that this is the case; indeed, there will typically be multiple suitable shadow values consistent with any finite dataset.

This model of taste change, which allows for adjustments in the willingness-to-pay for every good in every period, introduces as many free parameters as there are observations and so is clearly extremely permissive - even given its apparently restrictive additive form. Instead, suppose that we have prior grounds to believe that taste changes have been confined to a single good, denoted good 1. For example, in our empirical application, we suppose that good 1 represents tobacco products - a good for which there are strong grounds for believing that tastes have changed significantly given increased awareness of ill health effects. With the restriction of taste change to a single good, then $\boldsymbol{\alpha}_t = [\alpha_t^1, 0, ..., 0]'$, yielding the following temporal series of utility functions:

$$u(\mathbf{q}; \alpha_t^1) = v(\mathbf{q}) + \alpha_t^1 q^1 \tag{2}$$

Taste changes thus enter the basic utility maximisation framework in a more restricted manner. Specifically, the additive-linear specification for taste perturbations implies that the marginal rate of substitution between any of the other goods $j, k \in \{2, ..., K\}$ is invariant to taste instability on good-1. Another implication of this functional form is that preferences will obey the single crossing property.

[2] The concept of a virtual budget was first suggested by Rothbarth (1941) and Neary and Roberts (1980) to develop the theory of choice behaviour under rationing.

Definition (Milgrom and Shannon 1994) A utility function $u(\mathbf{q};\alpha_t^1)$ satisfies the single crossing property in $(\mathbf{q};\alpha^1)$ if for $\mathbf{q}' > \mathbf{q}''$ and $\alpha^{1\prime} > \alpha^{1\prime\prime}$

$$u(\mathbf{q}';\alpha^{1\prime\prime}) \geq u(\mathbf{q}'';\alpha^{1\prime\prime}) \Rightarrow u(\mathbf{q}';\alpha^{1\prime}) \geq u(\mathbf{q}'';\alpha^{1\prime}) \tag{3}$$

This condition can be interpreted as stating that for any $\mathbf{q}' > \mathbf{q}''$, the function $f(\alpha^1) = u(\mathbf{q}';\alpha^1) - u(\mathbf{q}'';\alpha^1)$ crosses the horizontal axis only once, from negative to positive, as α^1 increases. The single-crossing property means that there is an unambiguous ranking of tastes over time for the good of interest. For example, $\alpha_t > \alpha_s$ implies that a consumer's taste for good-1 is higher at t than at s. That is, the marginal rate of substitution with respect to good-1 is always higher at t than at s at every point in commodity space.[3]

2.1 Empirical conditions

We are interested in establishing whether a consumer's observed choice behaviour, D, could have been generated by taste change on a single good. Rationalisation of D by our model is defined as follows.

Definition 2 A consumer's choice behaviour, $D = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,\ldots,T}$, can be 'good-1 taste rationalised' by the base utility function $v(\mathbf{q})$ and a temporal series of additive linear perturbations to the marginal utility of good-1, $\{\alpha_t^1\}_{t=1,\ldots,T}$, if

$$v(\mathbf{q}) + \alpha_t^1 q^1 \leq v(\mathbf{q}_t) + \alpha_t^1 q_t^1 \tag{4}$$

for all $\mathbf{q}$ such that:

$$\mathbf{p}_t'\mathbf{q} \leq \mathbf{p}_t'\mathbf{q}_t \tag{5}$$

In words, the data can be rationalised by the model if there exists a time-invariant base utility function, and a series of perturbations to the marginal utility of good-1 such that observed choices are weakly preferred to all feasible alternatives. The empirical conditions, involving only observables, that are equivalent to a

[3] In the empirical section of this paper, we will motivate the use of quantile regression methods by appeal to the fact that an extension of our model, which directly includes interpersonal time-invariant heterogeneity, satisfies single crossing. In the two-good case, a basic result of single crossing is that $\arg\max_{q^1} u(q^1, q^0; \alpha^1)$ increases with α^1 (see Milgrom and Shannon (1994)).

rationalisation of the data by our theoretical model are given in Theorem 1.

Theorem 1. The following statements are equivalent:

1. Observed choice behaviour, $D = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$, can be good-1 taste rationalised.
2. There exist $3T$ real numbers $\{v_t, \alpha_t^1, \lambda_t > 0\}_{t=1,...,T}$, such that for all $s, t \in \{1, ..., T\}$

$$\begin{aligned} v_s - v_t + \alpha_t^1(q_s^1 - q_t^1) &\leq \lambda_t \mathbf{p}_t'(\mathbf{q}_s - \mathbf{q}_t) \\ \alpha_t^1 &\leq \lambda_t p_t^1 \end{aligned} \tag{6}$$

Proof: See Appendix A.

Theorem 1 is implied by optimising behaviour within the theoretical framework. If there exists a non-empty solution set to the inequalities defined by Theorem 1, then there exists a well-behaved base utility function and a series of taste shifters on good-1 that perfectly rationalise observed behaviour. The variables referred to by the revealed preference inequalities in part (2) of Theorem 1 have natural interpretations. The numbers $\{u_t\}_{t=1,..,T}$ and $\{\lambda_t\}_{t=1,..,T}$ can be interpreted respectively as measures of the level of baseline utility and the marginal utility of income at observed demands. The α_t^1 values can be interpreted as the marginal utility perturbation to good-1 relative to that dictated by base utility at observed demands since we can set $\alpha_t^1 = 0$ for all t.

Theorem 1 is an extension to the equivalence result originally derived by Afriat (1967) for the utility maximisation model with a time-invariant utility function. Imposing $\alpha_t^1 = 0$ for all $t = 1, .., T$ returns the standard Afriat inequalities. If there is no intertemporal variation in good-1, $q_t^1 = q_s^1$ for all $t, s \in 1, ..., T$, then we immediately have the standard Afriat conditions.

Given solutions for condition (2) in Theorem 1, we can construct a rationalising utility function at each observation $(u(\mathbf{q}_t; \alpha_t^1))$ and also examine counterfactuals such as at $u(\mathbf{q}_t; \alpha_s^1)$. This indicates the utility which would be derived from consuming the period t bundle, but with one's tastes from another period. For example, comparing $u(\mathbf{q}_t; \alpha_t^1)$ with $u(\mathbf{q}_t; \alpha_s^1)$ gives

$$u(\mathbf{q}_t; \alpha_t^1) - u(\mathbf{q}_t; \alpha_s^1) = \left(\alpha_t^1 - \alpha_s^1\right) q_t^1$$

whilst the shadow price of consumption of good 1 in period t given period s tastes is

$$u_1(\mathbf{q}_t;\alpha_s^1) = \lambda_t \left[p_t^1 - \frac{\left(\alpha_t^1 - \alpha_s^1\right)}{\lambda_t} \right]$$

Note that the shadow price of consumption in period t, given tastes in period s, may be negative if the taste-change term $\left(\alpha_t^1 - \alpha_s^1\right)$ is large enough. For example, if the consumer's taste for tobacco changes negatively between some earlier period $t-1$ and a later period t such that $\alpha_{t-1}^1 - \alpha_t^1$ is sufficiently positive then it is possible that

$$u_1(\mathbf{q}_{t-1};\alpha_t^1) = \lambda_{t-1} \left[p_{t-1}^1 - \frac{\left(\alpha_{t-1}^1 - \alpha_t^1\right)}{\lambda_{t-1}} \right] < 0$$

The interpretation of this is that the consumer in period t would need to be paid to smoke as much as they did back in period $t-1$.

Solutions to Theorem 1 also enable us to construct the virtual prices at which the individual with preferences given by the base utility function would have purchased the bundle of goods purchased at t with taste for tobacco α_t^1

$$\widetilde{p}_t^1 = p_t^1 - \frac{\alpha_t^1}{\lambda_t}$$

Interpreting taste change as an evolution of virtual prices supports the interpretation of information programmes as supplementary tax and incomes policies. For example, programmes designed to cultivate a negative taste for tobacco can be thought as levying a 'taste-tax' on the good because they manifest themselves in a rise in the virtual price for tobacco: $\tilde{p}_t^1 > p_t^1$ as $\alpha_t^1 < 0$ for a negative taste perturbation. Given the virtual price characterisation, variation in α^1 is more easily interpreted as a change in the marginal willingness to pay for good-1. The magnitude of the change in the marginal willingness to pay is captured by the term α_t^1/λ_t. This is useful because there is no clear behavioural interpretation of the magnitude of α_t^1 since its value depends upon the cardinal representation of base preferences.

2.2 Rationalisability

Surprisingly, under mild assumptions on D,[4] observed behaviour can *always* be explained by our simple model; that is, one can find sets of base-utility numbers, $\{v_t\}_{t=1,...T}$, marginal utilities of income $\{\lambda_t\}_{t=1,...,T}$ and taste perturbations on a single good $\{\alpha_t^1\}_{t=1,...,T}$ that rationalise observed choice behaviour.

Theorem 2 Given $q_t^1 \neq q_s^1$ for all t and s, any data set D can be good-1 taste rationalised.

Note that quantity variation is sufficient, but not necessary, for D to be good-1 taste rationalised. Let $\mathbf{p}_t^{\neg 1}$ denote the price vector for all goods except good-1, $\mathbf{p}_t^{\neg 1} = [p_t^2, ..., p_t^K]$, and define $\mathbf{q}_t^{\neg 1}$ analogously. For subsets $S_t \subseteq \{1, ..., T\}$ within which $q_t^1 = q^1$ for all $t \in S_t$, if the choice set $\{\mathbf{p}_t^{\neg 1}, \mathbf{q}_t^{\neg 1i}\}_{t \in S_t}$ satisfies the Generalised Axiom of Revealed Preference (GARP), then D will be rationalisable by our framework despite the violation of perfect variation in good-1. A leading example of this is when an agent does not buy good 1 in more than one period ($q_s^1 = q_t^1 = 0$ for some $s \neq t$).

Theorem 2 is closely related to Varian's (1988) Theorem 1, in which it is proved that the standard utility maximisation model is virtually emptied of empirical content if the price of at least one good is not observed. In such circumstances, one can hypothesize that the unobserved prices take on values high enough that expenditures on goods with unobserved prices dominate all other revealed preference comparisons. The virtual budget characterisation of taste change makes clear the connection between Theorem 2 and Varian's result: tastes for good-1 could always decline to the extent that the virtual prices required to support observed bundles are high enough to prevent an intersection of the virtual budget hyperplanes.

2.3 Recoverability

The revealed preference inequalities associated with our theoretical framework can be used to recover the set of *minimal* perturbations to the marginal utility of good-1 that will rationalise observed behaviour. This set is always non-empty if D satisfies perfect variation with respect to good-1. We here outline an easy to implement quadratic programming procedure that enables the recovery of minimal marginal utility pertur-

[4] Our assumption that we observe period-to-period variations in quantities such that $q_t^k \neq q_s^k$ for all $t, s = 1, ..., T$ is important here.

bations on good-1 that are sufficient to rationalise a data set D.

Recovery of α^1 The minimal squared perturbations to the marginal utility of good-1 relative to preferences in period 1 that are necessary to good-1 rationalise observed choice behaviour $D = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ are identified as the unique solution set $\{\alpha_t^1\}_{t=1,...,T}$ to the following quadratic programme.

$$\min_{\{v_t, \lambda_t, \alpha_t\}_{t=1,...,T}} \sum_{t=1}^{T} \left(\alpha_t^1\right)^2 \tag{7}$$

subject to:

1. The revealed preference inequalities:

$$\begin{aligned} v_s - v_t + \alpha_t^1(q_s^1 - q_t^1) &\leq \lambda_t \mathbf{p}_t'(\mathbf{q}_s - \mathbf{q}_t) \\ \alpha_t^1 &\leq \lambda_t p_t^1 \end{aligned} \tag{8}$$

2. The normalisation conditions:

$$\begin{aligned} v_1 &= \delta \quad \text{(an arbitrary constant)} \\ \lambda_1 &= \beta \quad \text{(a strictly positive constant)} \\ \alpha_1^1 &= 0 \end{aligned} \tag{9}$$

for $s, t = 1, ..., T$.

Minimising the sum of squared α^1's subject to the set of revealed preference inequalities ensures that the recovered pattern of minimal taste perturbations are sufficient to rationalise observed choice behaviour. The normalisation conditions are required because the quadratic programming problem is ill-posed in their absence. This is due to the invariance of tastes to positive monotonic changes in the utility function. Let $\{\bar{v}_t, \bar{\lambda}_t, \bar{\alpha}_t^1\}_{t=1,...,T}$ represent some feasible solution to the rationalisation constraints. As the set represents a feasible solution, the following inequalities hold for all $s, t \in \{1, ..., T\}$:

$$\bar{v}_s - \bar{v}_t + \bar{\alpha}_t(q_s^1 - q_t^1) \geq \bar{\lambda}_s \mathbf{p}_s'(\mathbf{q}_s - \mathbf{q}_t) \tag{10}$$

However, the following set of inequalities is also feasible:

$$\beta(\bar{v}_s + \delta) - \beta(\bar{v}_t + \delta) + \beta\bar{\alpha}_t^1(q_s^1 - q_t^1) \geq \beta\bar{\lambda}_s\mathbf{p}_s'(\mathbf{q}_s - \mathbf{q}_t) \tag{11}$$

for $\delta \in (-\infty, \infty)$ and $\beta > 0$. Thus, without a location and scale normalisation, there exist an infinite number of feasible sets of utility numbers that can be associated with a given set of feasible taste shifters. Without loss of generality, we impose $\alpha_1^1 = 0$, which allows us to interpret $\{\alpha_t^1\}_{t=2,...,T}$ as the minimal rationalising marginal utility perturbations to good-1 relative to preferences at $t = 1$. We also impose the scale normalisation $\lambda_1 = \beta$ to ensure that the output of the quadratic programming procedure is scaled sensibly.

3 Rationalising tobacco consumption

In this section we apply our revealed preference methodology to study taste changes with respect to tobacco in the UK. Over the period we examine (1980 to 2000), the UK population acquired more information about the ill-effects of smoking, perhaps causing tastes for tobacco to change. The Health and Lifestyle Survey 1984 and the Office for National Statistics Omnibus Survey 1996, for example, both asked questions about the connection between smoking and poor health. Between these years, there was a large rise in awareness of the link between smoking and heart disease; in 1984, only 25.8% of the sample believed that smoking caused heart disease, while by 1996, 80.6% recognised the link. Thus we argue that there is indeed strong *prima facie* evidence that tastes have changed in a particular direction, and we would therefore expect that our approach should be able to identify this.

3.1 The data

There is no UK consumption panel over this period that includes spending on tobacco and covers a sufficiently long time period with sufficient power to reject the null hypothesis of no taste change. However we do have rich cross section consumption data in the Family Expenditure Survey (FES), which records detailed expenditure and demographic information for 7,000 randomly selected households each year between 1961 and 2001.

From this dataset, we sample households whose head was aged between 25 and 35 years old in 1980 (i.e. they were born between 1945 and 1955).[5] between 1980 and 2000. We stratify by education level; the "low" education group, L, is formed from those individuals who left school no older than the legal minimum, 15 years old at the time. Those staying in school past this age comprise a "high education" group, H.

We are primarily concerned with choice over a tobacco aggregate and a nondurable commodity aggregate, i.e the number of goods is reduced to. $K = 2$. Appendix B provides a detailed list of goods that are classed as nondurables in our data set. Total expenditure is defined as all spending on these goods. Price indices are constructed for the tobacco and nondurable aggregates using the sub-indices of the U.K. Retail Price Index. As shown in Figure 1, the relative price of tobacco rose throughout much of the period except over the mid-eighties. Cigarette taxation in the UK consists of both a specific (per-pack) and an ad-valorem component (as a percentage of the retail pack price). Other tobacco products such as cigars and hand-rolled tobacco are subject to their own specific duties. In addition to these, VAT is charged on the pack price (inclusive of the specific and ad-valorem components of the duty). The real value of total duty (including both the specific and ad-valorem components) charged on a packet of 20 cigarettes has grown quite considerably over the period covered by our data *except* for a period in second half of the 1980's when excise duties as a percentage of the retail price in fact fell.[6] An explicit tobacco duty was introduced in 1993 and between 1993 and 2000 when our data ends, the specific component of duty was increased initially by 3% above inflation (as measured by the Retail Prices Index (RPI)) and then 5% from July 1997. In 2001, at the end of the period covered by our data the government stopped the escalator and froze duties in real terms.

In the following sections, we will rationalise the changes in tobacco consumption at six quantiles of the budget share for tobacco distribution conditional on education. Figure 2 displays the (unconditional) budget shares for these quantiles between 1980 and 2000. At each quantile, the highly education individuals devotes a smaller proportion of their budget to tobacco than the low education individuals. In fact, for the high education group, the demand of the 0.65 quantile individual falls to zero in 1983 and remains nil for the the rest of the period considered, while the 0.75 quantile is zero by the end of the period. Tobacco consumption

[5] We select this birth cohort because less than 5% of smokers start smoking after they reach their 25^{th} birthday (Office for National Statistics, 2012). The assumption that the population of smokers is stable is then relatively mild.

[6] See Level, O'Connell and Smith, Ch 9, *The Green Budget 2016*, Figure 9.5.

Figure 1: Price of Tobacco Relative to Nondurables, 1980-2000

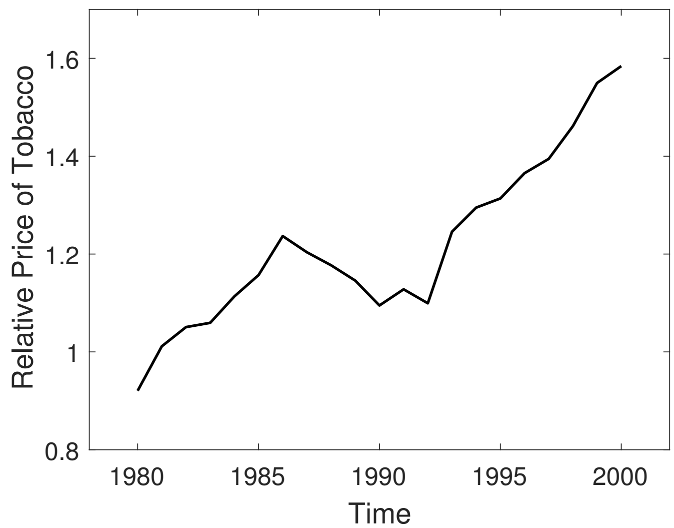

remains strictly positive in all periods for the quantiles of the low education individuals that are considered.

Figure 2: Quantiles 0.65 to 0.90 of the Budget Share of Tobacco Distribution, 1980-2000

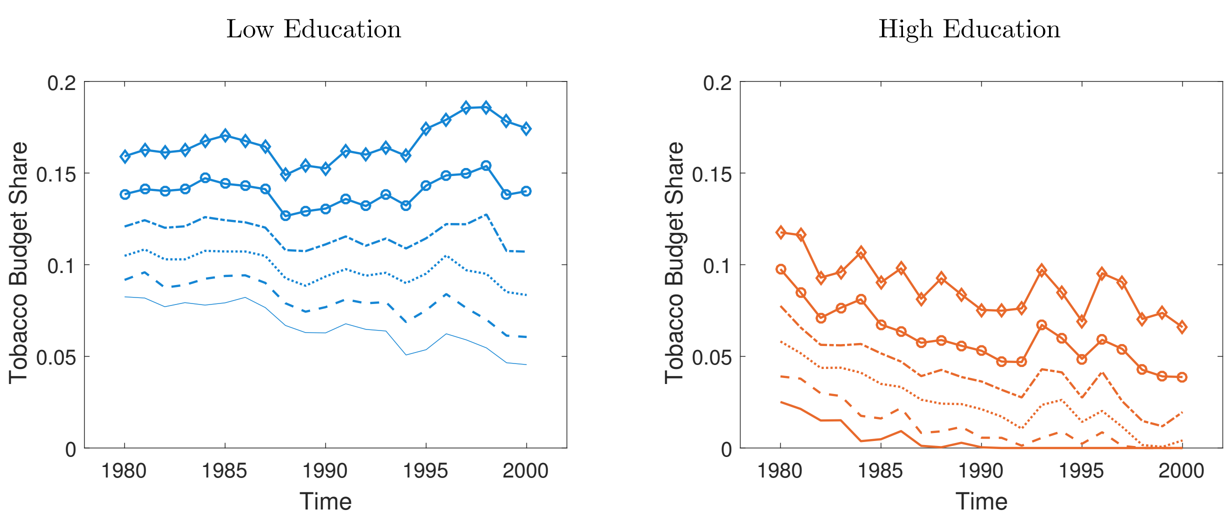

3.2 Quantile demands

In experimental revealed preference analyses, recovery exercises are typically conducted on an individual-by-individual basis in the laboratory. Indeed, this is the implicit setting assumed in our theoretical set-up. However, our aim here is to apply the analysis developed above for a long time-series of repeated cross-section consumer expenditure data. Given the structure of our data, we develop conditions that ensure quantile demands recover individual demand behaviour over time. We then apply the revealed preference methodology to estimated quantities in order to make statements about taste changes over time. This approach places restrictions on unobserved time invariant interpersonal heterogeneity which we now address.

We observe individuals $i = 1, ..., N_t^c$ for each education level $c = \{L, H\}$ in each period $t = 1, ..., T$. For notational simplicity, the education and individual subscripts are suppressed for the rest of this section. Drawing on Blundell, Kristensen, and Matzkin (2014), we restrict the dimensionality of the demand system to $K = 2$ and assume that individual demands are monotonic in scalar unobserved heterogeneity. Good 1 is tobacco, our good of interest. To allow for time invariant interpersonal heterogeneity, we augment individual preferences in (2) to take the form:

$$u(\mathbf{q}; \alpha_t) \quad = \quad f(q^1, q^0, \tau) + \alpha_{t,\tau}^1 q^1 \tag{12}$$

where $\tau \sim U(0, 1)$ represents time-invariant interpersonal taste heterogeneity and $\alpha_{t,\tau}^1$ gives the perturbation to the marginal utility of good 1 experienced by that individual consumer at time t, as in (2).

For quantile demands to identify individual demands we have to place restrictions on preferences that guarantee demands are monotonic in unobserved scalar heterogeneity. Matzkin (2007) provides the following general conditions on heterogeneous preferences $f(q^1, q^0, \tau)$ to be invertible in unobserved heterogeneity:

$$f(q^1, q^0, \tau) \quad = \quad f(q^1, q^0) + w(q^1, \tau) \tag{13}$$

where the functions f and w are twice continuously differentiable, strictly increasing and strictly concave, and $w(q^1, \tau)$ has a strictly positive second derivative. Under these conditions, the demand function for q^1 will satisfy the restrictions of consumer choice for each value τ. Budget shares will also be monotonic in τ.

Preferences displaying taste change are the given as

$$u(\mathbf{q}; \alpha_t) \quad = \quad f(q^1, q^0) + w(q^1, \tau) + \alpha_{t,\tau}^1 q^1, \tag{14}$$

which, without loss of generality can be represented as:

$$u(\mathbf{q};\alpha_t) \quad = \quad f(q^1, x - p^1 q^1) + \tau q^1 + \alpha^1_{t,\tau} q^1. \tag{15}$$

With only interpersonal time-invariant heterogeneity (i.e. $\alpha^1_{t,\tau} = 0$ for all $t = 1, ..., T$), invertibility guarantees the family of preferences $\{f(q^1, x - p^1 q^1) + \tau q^1\}$ obeys single-crossing. Thus, in this context, quantile demands, conditional on $(\mathbf{p}, x)$, identify individual demands (see Blundell et al (2014) for further discussion). Let $q^1(\mathbf{p}, x, \tau)$ give the demand for good 1 that is associated with taste heterogeneity τ. Then,

$$\begin{aligned} q^1(\mathbf{p}, x, \tau) &= \max_{q^1} f(q^1, x - p^1 q^1) + \tau q^1 \\ &= Q_{q^1}(\tau|\mathbf{p}, x) \end{aligned} \tag{16}$$

where $Q_{q^1}(\tau|\mathbf{p}, x)$ denotes the τ'th conditional quantile of q^1.

To equate quantile and individual demands given the time varying component to unobserved preference heterogeneity, $\alpha^1_{t,\tau}$, we further assume that taste evolution term $(\tau + \alpha^1_{t,\tau})q^1$ does not alter the ranking of consumers over time in the q^1 distribution conditional on $(\mathbf{p}, x)$. With this assumption, the τ'th conditional quantile reflects the demand behaviour of an individual with unobserved time-invariant preference heterogeneity τ and $\alpha^1_{t,\tau} q^1$.

While a strong assumption, we do not believe rank invariance to be incredible in this context; we only consider individuals who are old enough to have made an initial decision to smoke (see above) and high profile anti-smoking campaigns are not targeted in a way that would cause one to expect significant re-ranking of individuals within education groups. More concretely, a testable implication of rank invariance is covariate similarity at the same quantile of the tobacco distribution over time. Table 1 reports the results from Dong and Shen's (2016) distributional tests for rank similarity at the quantiles we follow in our primary results between the start and end period of our analysis. We cannot reject the null hypothesis of rank invariance on the basis of the available data at conventional significance levels.

Table 1: Tests for Rank Invariance of Conditional Quantiles by Education Level

	High Education		Low Education	
Quantile	χ^2		χ^2	
0.65	5.929	(0.655)	10.980	(0.203)
0.70	8.219	(0.412)	12.006	(0.151)
0.75	5.131	(0.744)	12.874	(0.116)
0.80	6.327	(0.616)	13.133	(0.107)
0.85	8.372	(0.398)	9.073	(0.336)
0.90	9.989	(0.266)	8.439	(0.392)

Notes: Tests based on Dong and Shen (2015). Covariance matrix bootstrapped with 1,000 replications. P-values in parentheses. Quantiles estimated conditional on household composition, expenditure and expenditure squared. Covariate similarity is judged with respect to sex, age and employment status.

3.3 Censoring

We proceed by estimating conditional quantile demands in terms of budget shares. As is clear from Figure 1, the latent budget share for tobacco, $w^\star$ is left censored at zero: $w = max\{0, w^\star\}$. To allow for censoring, we follow the estimation method of Chernozhukov and Hong (2002) and Chernozhukov, Fernández-Val, and Kowalski (2015).

For each of the $t = 1, ..., T$ price regimes, we consider the following censored quantile equations for the budget share of tobacco (for ease of notation, cohort and time subscripts are suppressed):

$$\begin{aligned} w &= max(0, w^\star) \\ w^\star &= Q_{w^\star}(\tau | x, z) \end{aligned} \tag{17}$$

where without loss of generality

$$\tau \sim U(0,1)|x, \mathbf{z} \tag{18}$$

and x is total expenditure (on nondurables and tobacco) and $\mathbf{z}$ is a vector of household characteristics.

To allow for the possible endogeneity of total expenditure x we follow Blundell and Powell (2007) and Imbens and Newey (2009) to use a quantile control function approach. We use disposable income as an excluded instrument that allows us to recover the control variable. The estimator we adopt for this quantile control function model, introduced byChernozhukov, Fernández-Val, and Kowalski (2015), is given by

$$Q_{w^\star}(\tau | x, \mathbf{z}, v) = \mathbf{X}'\boldsymbol{\beta}(\tau) \tag{19}$$

where v is the unobserved latent control variable that is included to account for the possible endogeneity of x. In implementation, $\mathbf{X}$ is replaced with $\hat{\mathbf{X}} = [log(x), log(x)^2, z, \widehat{v}]$, where $\widehat{v}$ is the conditional quantile regression estimate of v and where the the OECD demographic index is used to capture demographic characteristics, z.

For each education level at each price regime, $\boldsymbol{\beta}(\tau)$ is estimated as:

$$\boldsymbol{\beta}(\tau) = \arg\min_{\boldsymbol{\beta}} \sum_{i \in S} \rho_\tau(w_i - \hat{\mathbf{X}}_i' \boldsymbol{\beta}) \tag{20}$$

where ρ_τ is the standard asymmetric absolute loss function of Koenker and Bassett Jr (1978) and S denotes the set of observations for which $Pr(w_i > 0) > \delta$, i.e. those for which the probability of censoring is negligible and a linear functional form for the conditional quantile is justified. We refer the reader to Chernozhukov, Fernandez-Val and Kowalski (2015) for a more detailed exposition, including the construction of S. We then evaluate demands in each period at a set of expenditure levels set to ensure that budget lines at different time periods intersect. This is a necessary condition for detecting violations of a time-invariant utility function in the data as we discuss in the next sub-section.

3.4 Power

We are interested in recovering the contribution of taste change to the trends in the distribution of tobacco consumption over time and between different quantiles. As previously discussed in the literature, budget variation in the data hinders our progress somewhat because if budget constraints do not cross then taste changes will be undetectable. A second issue, not yet been widely acknowledged in the literature on detecting violations of revealed preference, is that in a two-good demand system the direction of price changes matters for our ability to recover the systematic taste change over time. We here discuss the modifications made to our approach that are intended to make our methodology as powerful as possible.

Income Growth Income growth in survey data prevents the intersection of budget hyperplanes, weakening ones ability to detect violations of stable preferences (see Varian (1983)). We thus evaluate each quantile's demands along a Sequential Maximum Power (SMP) path (Blundell, Browning, and Crawford 2003). Intuitively, the SMP method controls for budget variation to provide a choice set that maximises the chance of detecting violations of a time-invariant utility function in observational data. Along an SMP path, expenditure levels are set such that demands at $t+1$ are just weakly preferred to demands at t:[7]

$$\widetilde{x}_t = p_t' q_{t-1}^{SMP} \tag{21}$$

We can abstract from the complicating issues caused by transitivity in the construction of the SMP path because, as first highlighted by Rose (1958), transitivity has no empirical content for a 2-good demand system (Blundell, Browning, Cherchye, Crawford, De Rock, and Vermeulen 2015).

Price changes Qualitative evidence suggests that tastes for tobacco decreased over the period of interest. Given our two-good framework, we will then only reject a time-invariant utility function over periods with a falling relative price of tobacco. The construction of the SMP path imposes $p_{t+1}' q_{t+1}^{SMP} = p_{t+1}' q_t^{SMP}$ between adjacent periods. A violation of revealed preference between these periods requires $p_t' q_t^{SMP} > p_t' q_{t+1}^{SMP}$. Combining these conditions implies that we require

$$\left[\frac{p_{t+1}^1}{p_{t+1}^0} - \frac{p_t^1}{p_t^0}\right](q_{t+1}^1 - q_t^1) > 0$$

to detect negative taste changes for tobacco; i.e. violations of GARP only occur when the relative price and the quantity changes move in the same direction in our two-good SMP framework. With declining taste change, quantity demanded will always fall with increases in price; it is as if such consumers are very sensitive to price changes. Our method will thus be able to recover falls for tastes in tobacco over the seven-year period 1986-1992 in which relative prices for tobacco fell in the UK. To illustrate, Figure 3 gives the relative price of tobacco and unconditional budget shares for the tobacco quantiles of the high education group.

[7] In the first period, we 'start' the SMP path at the mean of nondurable expenditures of the low-education group. Expenditure levels are then adjusted such that the demand of the .75-quantile (our 'middle' quantile) is weakly affordable and evaluate demands of all quantiles and education groups at the same set of expenditure levels.

Figure 3: Falling Relative Prices & Budget Shares for Tobacco (High Education Group)

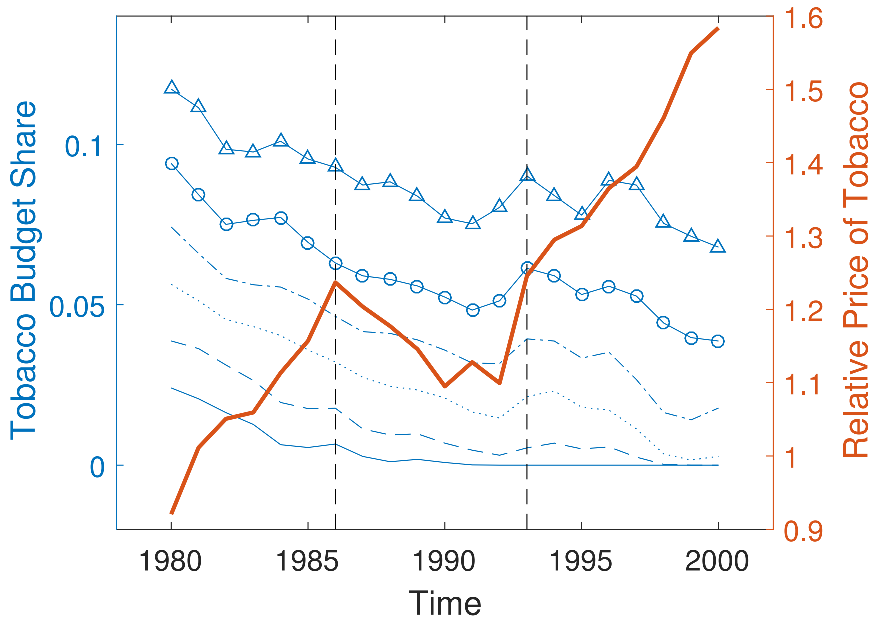

4 Results

Our revealed preference method recovers significant differences in tastes for tobacco across low and high education individuals. Figure 4 gives the minimum perturbations to the marginal utility of tobacco that are necessary to rationalise the consumption behaviour of the low and high-education individuals, given the normalisation of taste shifters relative to the heaviest-smoking (0.9 quantile) low-education individual in 1986. These perturbations are everywhere negative with the exception of tastes for the heaviest smoking, low-education individual in 1987. There is also a negative gradient to the taste trajectories we uncover, suggesting that tastes for tobacco declined over the period. Table 2 gives the 95% confidence interval on these marginal utility perturbations, showing that there is never a statistically significant difference in taste for tobacco for the heaviest smoking, low-education individual from their first-period tastes. Thus, tastes for the heaviest smoking, low-education individual do not appear to have changed over the period of interest.

Table 3 gives the educational differences in tastes across the tobacco distribution with the bootstrapped 95% confidence interval. We estimate the minimum marginal utility perturbation necessary to rationalise each bootstrapped quantity sequence, and within each set of rationalising perturbations calculate the dif-

Table 2: Marginal Utility Perturbations with 95% Confidence Interval, 1986-1992

Year	*Quantile*					
	.65	.70	.75	.80	.85	.90
Low Education						
1986	-6.4696	-5.4696	-4.9696	-3.9696	-2.0801	0
	(-13.11,-4.3)	(-12.46,-3.57)	(-11.41,-2.86)	(-10.41,-2.09)	(-9.22,-0.79)	-
1987	-6.946	-6.446	-5.946	-4.8611	-1.7512	0.5839
	(-13.99,-4.79)	(-13.23,-4)	(-11.92,-3.15)	(-10.68,-2.3)	(-9.57,-0.5)	(-1.16,1.69)
1988	-7.1759	-6.6759	-6.1759	-5.1759	-3.4034	-1.8621
	(-14.22,-5.31)	(-13.47,-4.59)	(-12.31,-3.81)	(-11.51,-3.1)	(-10.17,-2.03)	(-2.63,0.08)
1989	-7.7802	-7.2095	-6.2095	-5.7095	-3.3859	-0.2148
	(-14.51,-5.73)	(-13.47,-4.9)	(-12.63,-3.93)	(-11.44,-3.02)	(-10.36,-1.71)	(-1.92,0.9)
1990	-8.0882	-6.7746	-6.1924	-4.7293	-3.3949	-0.7264
	(-14.08,-5.52)	(-13.29,-4.66)	(-12.43,-3.83)	(-11.38,-2.87)	(-10,-1.49)	(-2.07,0.73)
1991	-7.8935	-6.8935	-6.3206	-5.3206	-3.3784	-0.493
	(-14.82,-5.74)	(-13.57,-4.73)	(-12.61,-3.81)	(-10.91,-2.91)	(-10,-1.55)	(-2.1,0.49)
1992	-9.117	-8.117	-7.2107	-6.2107	-4.166	-1.5072
	(-16.09,-7.02)	(-14.71,-6.02)	(-13.73,-4.98)	(-12.29,-3.86)	(-10.45,-2.3)	(-3.68,0.39)
High Education						
1986	-10.2225	-8.6492	-7.4696	-6.9696	-5.9696	-4.4696
	(-17.12,-7.62)	(-15.85,-6.42)	(-14.66,-5.45)	(-13.63,-4.56)	(-12.07,-3.08)	(-10.48,-1.48)
1987	-10.4898	-9.7974	-9.0027	-7.946	-7.446	-5.446
	(-17.44,-8.44)	(-17.06,-7.69)	(-16.06,-6.52)	(-14.81,-5.73)	(-13.59,-4.72)	(-11.71,-3.18)
1988	-10.5929	-9.9431	-9.2866	-8.1759	-7.6759	-5.6759
	(-17.5,-8.72)	(-16.99,-7.85)	(-15.97,-6.87)	(-14.98,-5.95)	(-13.71,-4.87)	(-11.58,-2.86)
1989	-10.7274	-9.973	-8.7802	-8.2802	-6.7095	-5.2095
	(-17.55,-8.54)	(-16.82,-7.76)	(-15.88,-6.92)	(-14.97,-5.89)	(-13.14,-4.32)	(-10.95,-2.6)
1990	-10.8859	-10.3859	-9.8859	-8.975	-7.2746	-5.6515
	(-17.92,-9.09)	(-17.53,-8.52)	(-16.7,-7.77)	(-15.64,-6.18)	(-13.87,-4.78)	(-11.74,-3.29)
1991	-10.7274	-9.812	-9.0363	-8.3935	-7.3935	-5.8206
	(-17.71,-8.59)	(-17.14,-8.06)	(-16.28,-7.25)	(-15.16,-6.33)	(-13.96,-5.14)	(-11.88,-3.68)
1992	-11.117	-10.617	-10.117	-9.617	-8.617	-6.7107
	(-18.39,-9.29)	(-17.83,-8.89)	(-17.22,-8.34)	(-16.38,-7.53)	(-14.92,-6.17)	(-12.93,-4.3)

Notes: Table gives the marginal utility of tobacco relative to the .9 quantile individual of the low-education group in 1986. We bootstrap the quadratic programming procedure, drawing 1,000 random samples with replacement from each education-year cell and report the associated 95% confidence interval on these perturbations.

Figure 4: Marginal Utility Perturbations - Quantiles 0.65 to 0.9 of the Budget Share of Tobacco Distribution, 1986-1992

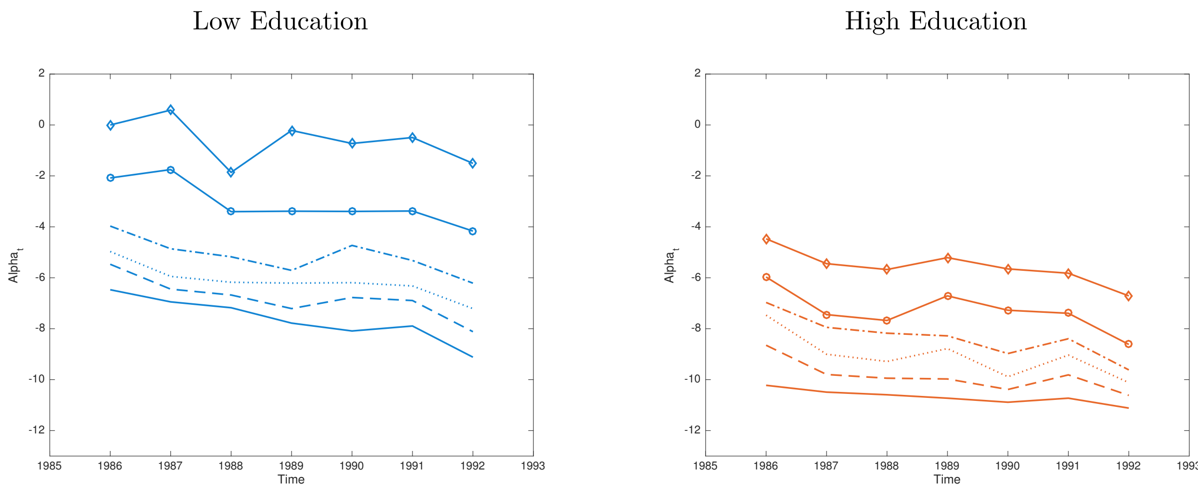

Notes: Figure 4 gives the minimum perturbations to the marginal utility of tobacco, $\alpha^1_{te\tau}$ necessary to rationalise estimated quantity sequences. Tastes are normalised relative to the .9 quantile individual of the low-education group in 1986 $\alpha^1_{1986,L,.9} = 0$. Lines closer to zero represent higher quantiles.

ference in low and high education tastes at each quantile. This exercise confirms that tastes for tobacco amongst the high-education group are significantly lower than the low-education group at all points in the distribution. There is also evidence that tastes amongst the heavy smokers are diverging, and that tastes amongst the light smokers are converging, across education groups over time. While taste differences between the low and high-education group decline between 1986 and 1992 for the .65 and .7- quantile, Table 3 suggests that this difference is increasing at the heavier smoking quantiles.

Table 3: Taste Differences Across Education Groups with 95% Confidence Interval, 1986-1992

	Quantile					
Year	.65	.70	.75	.80	.85	.90
1986	-3.7529	-3.1796	-2.5000	-3.000	-3.8896	-4.4696
	(-5.18,-2.00)	(-5.03,-2.00)	(-4.02,-2.00)	(-3.66,-2.00)	(-4.93,-1.01)	(-10.48,-1.48)
1987	-3.5438	-3.3514	-3.0568	-3.0849	-5.6948	-6.0299
	(-4.88,-2.31)	(-5.06,-2.47)	(-4.74,-2.50)	(-4.38,-2.50)	(-6.47,-2.50)	(-12.09,-3.14)
1988	-3.4171	-3.2672	-3.1107	-3.000	-4.2725	-3.8137
	(-4.49,-2.36)	(-4.57,-2.27)	(-4.04,-2.14)	(-3.47,-2.00)	(-4.59,-1.50)	(-10.91,-1.75)
1989	-2.9473	-2.7635	-2.5706	-2.5706	-3.3237	-4.9947
	(-3.97,-1.67)	(-4.11,-1.99)	(-4.15,-2.00)	(-3.99,-1.56)	(-5.13,-1.03)	(-11.05,-2.07)
1990	-2.7976	-3.6113	-3.6935	-4.2457	-3.8797	-4.9251
	(-4.48,-2.29)	(-4.8,-2.62)	(-5.04,-2.79)	(-5,-2.17)	(-5.86,-2.00)	(-11.12,-2.48)
1991	-2.8339	-2.9185	-2.7157	-3.0729	-4.0151	-5.3276
	(-3.92,-1.70)	(-4.46,-2.01)	(-4.58,-2.50)	(-4.33,-2.50)	(-5.67,-2.19)	(-11.55,-2.76)
1992	-2.000	-2.5000	-2.9063	-3.4063	-4.451	-5.2034
	(-3.30,-1.00)	(-3.91,-1.50)	(-4.45,-1.95)	(-4.82,-2.00)	(-6.23,-2.17)	(-12.57,-2.02)

Notes: Table gives the difference in the marginal utility of tobacco for low and high education groups, $\alpha^H_{\tau,t} - \alpha^L_{\tau,t}$. We bootstrap the quadratic programming procedure, drawing 1,000 random samples with replacement from each education-year cell and report the associated 95% confidence interval.

The negative gradient of the taste change trajectories depicted in Figure 4 suggests that tastes for tobacco declined over the period of interest. Figure 5 confirms this. For the high education individuals, tastes for tobacco in 1992 were statistically significantly lower than tastes in 1986 at the .75 to .90 quantiles. We cannot reject the null of stables tastes at the 95% confidence interval for the lightest smoking individuals (those at the .65 and .70 quantiles). However, this is not surprising given that tobacco consumption for these individuals was already very low in 1986, and we do not observe the price variation required for rejections prior to this date in our dataset. Thus, it appears that these groups had already incorporated the negative health information into their preferences before the period in which we are able to detect declining taste change. Tastes are converging overtime for the lighter smoking individuals as the falls in the marginal utility of tobacco for those with low education outstripped those of the high-education group, whose tastes were static over the period.

To put the value of the marginal utility perturbations in context, Figure 7 gives the implied virtual prices for tobacco, $\widetilde{p}^1$:

$$\widetilde{p}^1_{t,e,\tau} = p^1_t - \frac{\alpha^1_{t,e,\tau}}{\lambda_{t,e,\tau}} \tag{22}$$

These represent the change in the marginal willingness to pay for tobacco relative to base tastes along the SMP expenditure paths. As taste change has been negative, virtual prices will be higher than observed relative prices. While SMP expenditure paths do much to reduce the growth in expenditure, they do not eliminate it. Thus, declining marginal utility of income (λ), can act to magnify the declines in tastes.

Unsurprisingly, the virtual price of heavier smokers is lower compared to that of lighter smokers, and the virtual prices of the low-education group are lower than the high-education group. While the relative price of tobacco fell over the period of interest, virtual prices for many of our 'individuals' rose. This reflects the fact that declining tastes can be considered equivalent to a 'taste tax' was levied on tobacco in our framework. The exceptions are the individuals for whom we saw little change in tastes: the heaviest smoking individual with low-education and the light smoker with high-education. The virtual price trajectories had a negative gradient for these quantiles, although the level of virtual prices were much higher for the .65 quantile individual in the high-education group.

Figure 5: Taste Differences 1992-1986 with 95% Confidence Intervals - Quantiles 0.65 to 0.9 of the Budget Share of Tobacco Distribution

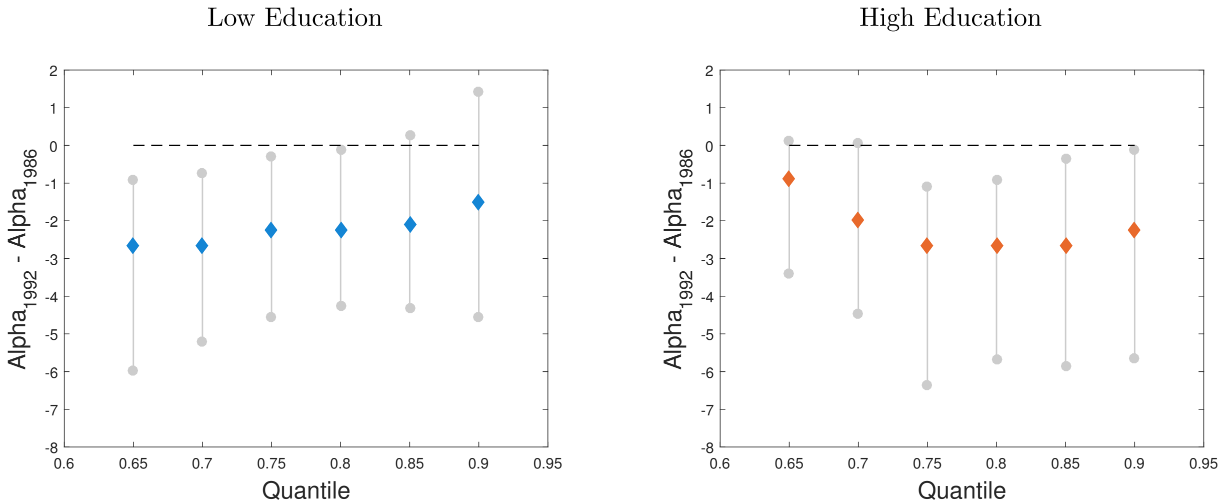

Notes: Figure 5 gives the difference in the marginal utility of tobacco between 1992 and 1986. Negative values indicate a decline in tastes over the period. Bootstrapped 95% confidence intervals are given. We draw 1,000 random samples with replacement from each education-year cell, run our quadratic programming procedure on each of the 1,000 quantity sequences and calculate the difference in beginning and end tastes for each quantity sequence.

Figure 6: Taste Wedge - Quantiles 0.65 to 0.9 of the Budget Share of Tobacco Distribution

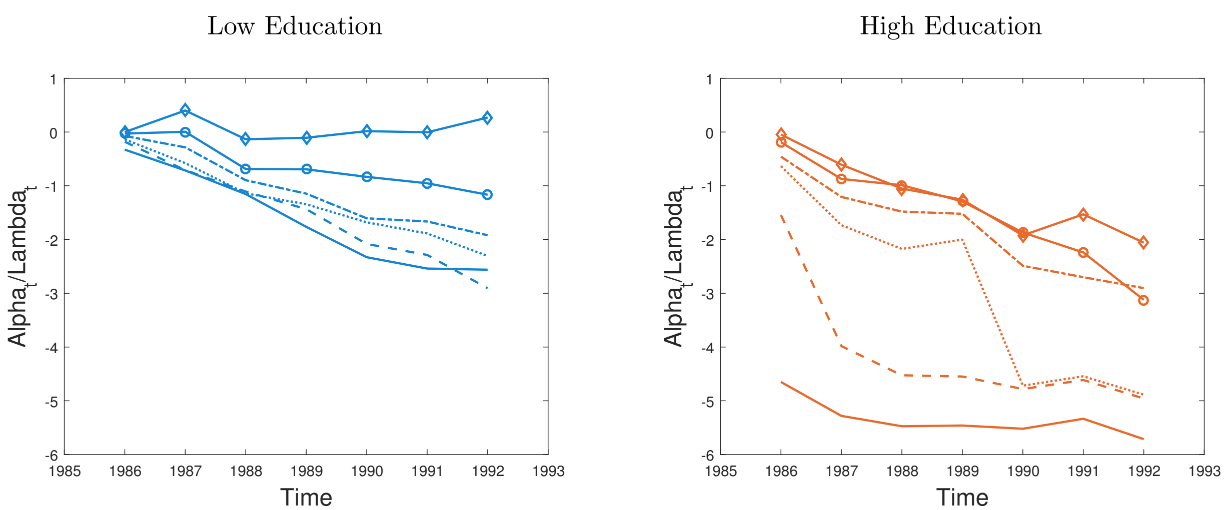

Notes: Figure 6 gives the minimum 'taste wedge', $\alpha^1_{t,e,\tau}/\lambda_{t,e,\tau}$, required to rationalise estimated quantity sequences. Tastes are normalised relative to the low-education .9 quantile in 1986.

Figure 7: Virtual Prices - Quantiles 0.65 to 0.9 of the Budget Share of Tobacco Distribution, 1986-1992

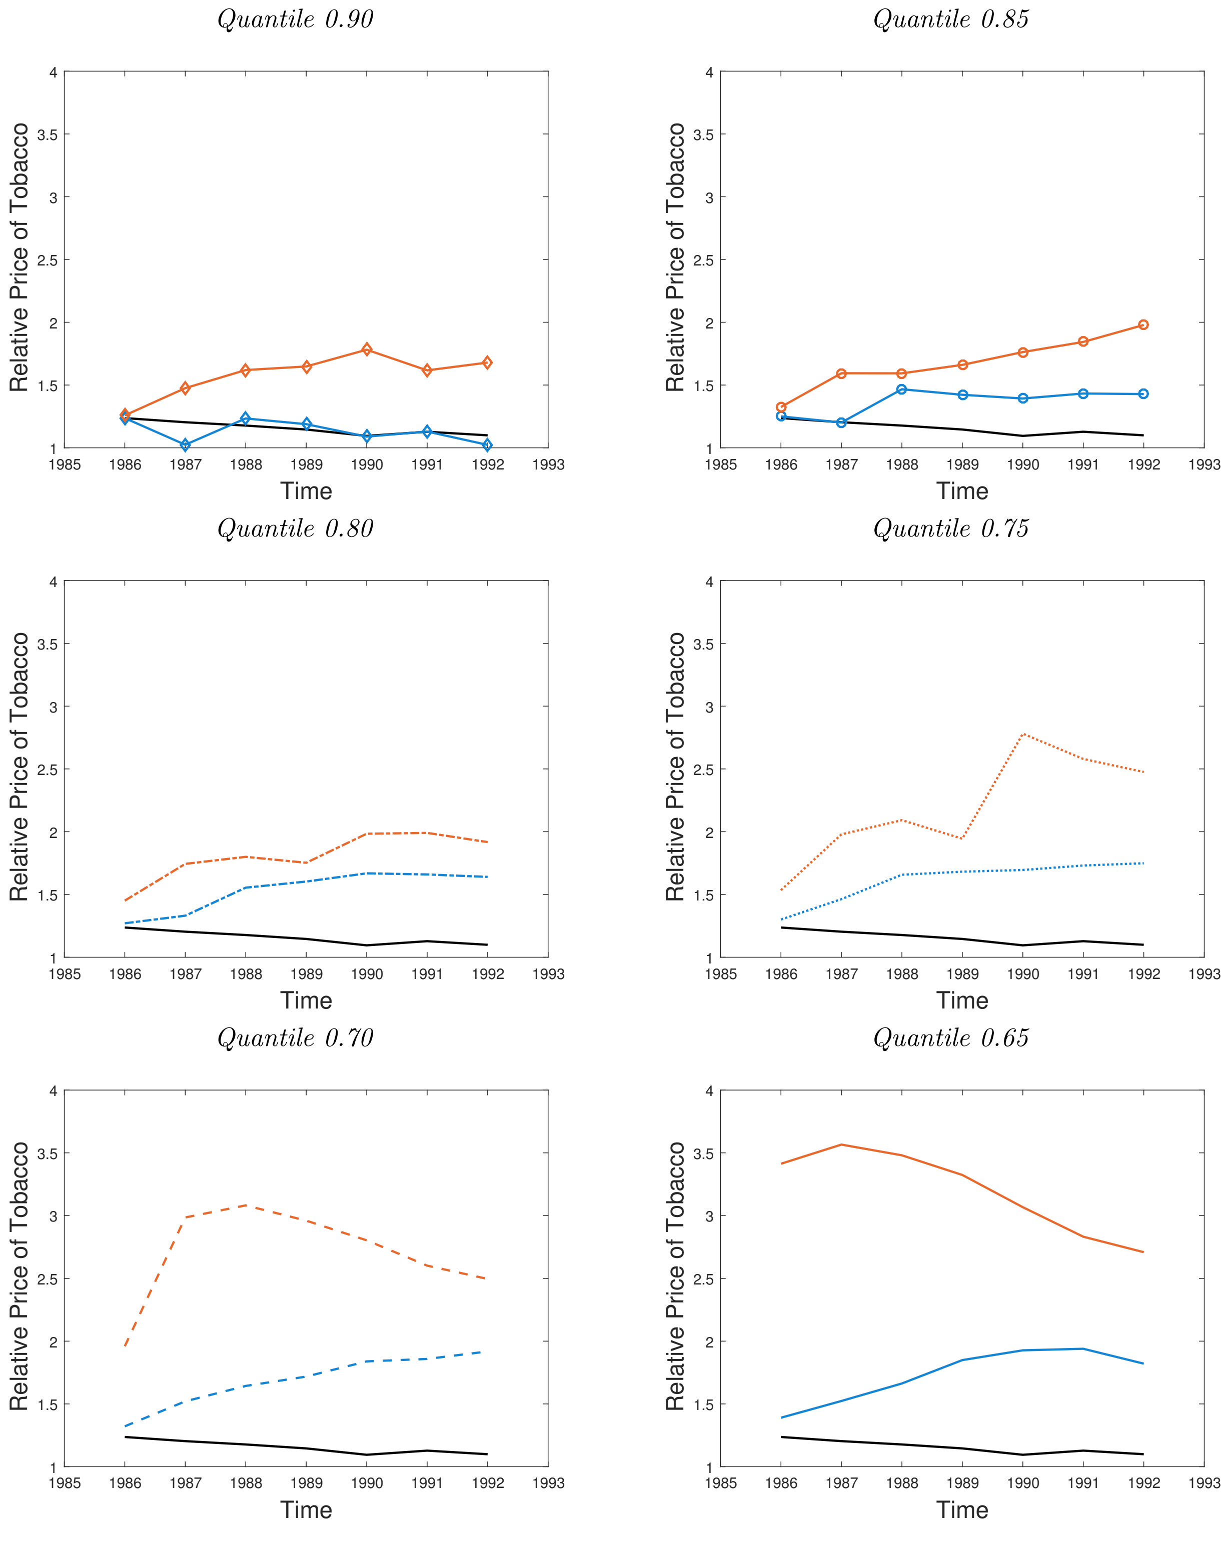

Notes: Figure 7 gives the virtual price of tobacco. These are the price sequences required to induce an individual with the tastes of the heavy-drinking, .9-quantile individual in 1986 to choose the estimated quantity sequences.

5 Conditional Demands and Taste Change

The estimation strategy we used above implicitly assumes that tobacco is weakly separable from all other non-durables, including alcohol. To examine the robustness of our findings to this assumption, and to explore whether additional patterns emerge from the data once this condition is relaxed, we follow the approach of Browning and Meghir (1991) and re-run our quadratic programming procedure on quantile demands that are estimated *conditional* on alcohol consumption. Specifically, we partition the set of observations into "light" and "heavy" drinkers depending on whether an individual consumes below or above the median budget share for alcohol in the first period. Given sample size limitations, this procedure is only implemented for low-education individuals. The quantile regression method outlined in the previous section is then applied to estimate demands within each alcohol-time cell at the expenditures along the same SMP path that was calculated previously. Base preferences are normalised relative to those of the heaviest smoking, heavy drinking, low education individuals in 1986. Please note that the fact that high-education demands are not included at this stage implies that the magnitude of estimated taste shifters are not comparable to the unconditional quantile results of the previous section.

We find our low drinking individuals have typically experienced more taste change than heavy drinkers in each smoking level-drinking level cell. Figure 9 shows that there were no statistically significant changes in taste for heavy-drinkers; the confidence interval on changes in tastes between 1986 and 1992 contains zero at all quantiles of the tobacco distribution. This finding is consistent with government and health practitioner reports that note low smoking cessation rates amongst heavy drinkers (Dollar, Homish, Kozlowski, and Leonard 2009). It also motivates the use of further restrictions on the joint consumption of alcohol and tobacco (e.g. bans on smoking in pubs and bars). Such restrictions would exploit the complementarity between alcohol and tobacco to reduce smoking amongst the heavy-smoking heavy-drinking group given the persistence in their tastes for tobacco.[8]

[8]See Adda and Cornaglia (2010) for additional effects of such restrictions.

Figure 8: Taste Change - Quantiles 0.65 to 0.9 of the Budget Share of Tobacco Distribution by Alcohol Consumption (Low Education), 1986-1992

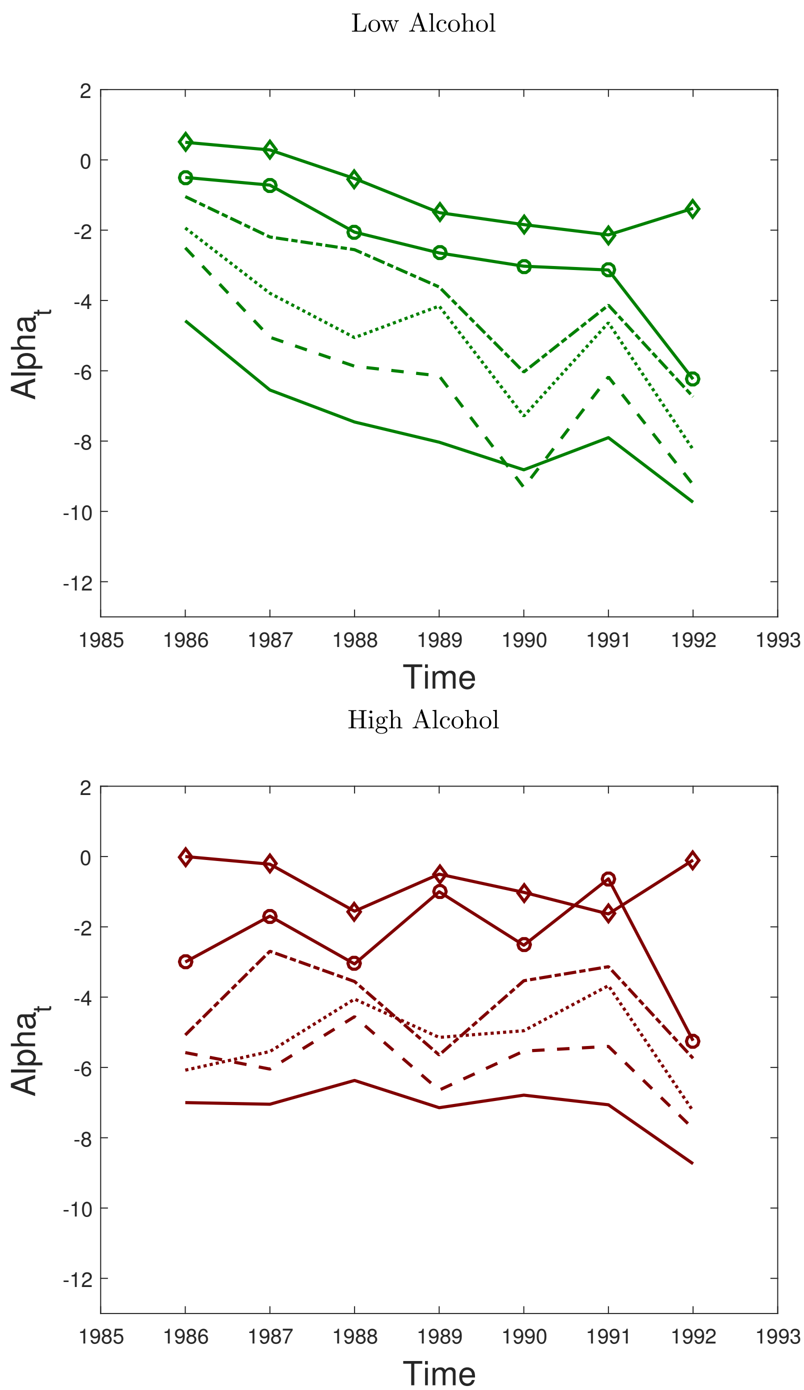

Notes: Figure 8 gives the minimum perturbations to the marginal utility of tobacco, $\alpha^1_{t,a,\tau}$ necessary to rationalise estimated quantity sequences. Tastes are normalised relative to the .9 quantile individual of the heavy drinking group in 1986 $\alpha^1_{1986,H,.9} = 0$. Lines closer to zero represent higher quantiles.

Figure 9: Difference in Tastes 1986-1992 - Quantiles 0.65 to 0.9 of the Budget Share of Tobacco Distribution by Alcohol Consumption (Low Education)

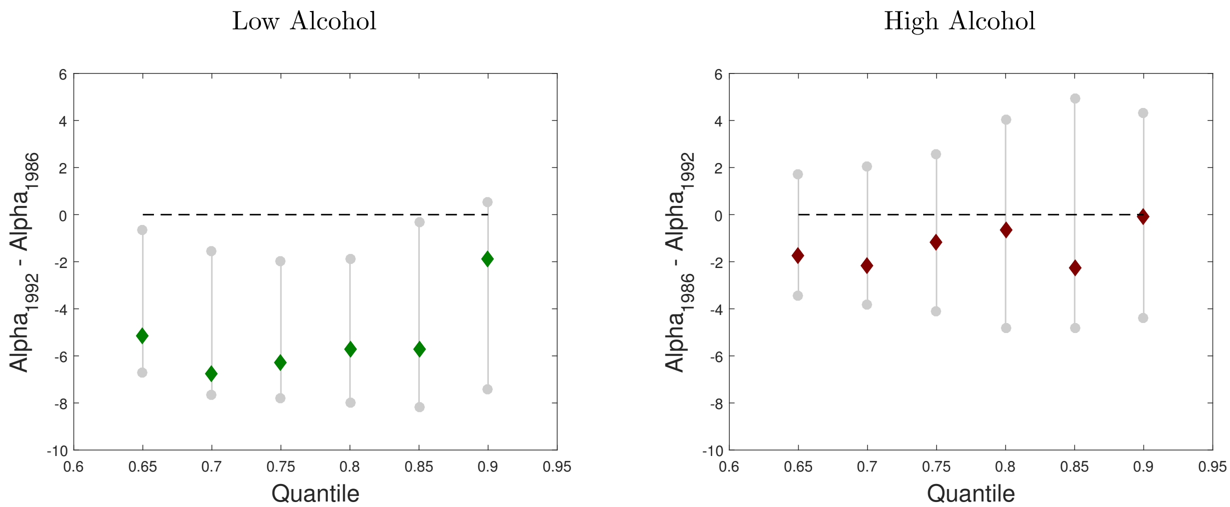

Notes: Figure 9 gives the difference in the marginal utility of tobacco between 1992 and 1986. Negative values indicate a decline in tastes over the period. Bootstrapped 95% confidence intervals are given. We draw 1,000 random samples with replacement, run our quadratic programming procedure on each of the 1,000 quantity sequences and calculate the difference in beginning and end tastes for each quantity sequence.

6 Conclusions

This paper has provided a theoretical and empirical framework for characterising taste change. We have uncovered a surprising non-identification result: observational data sets on a K-dimensional demand system can always be rationalised by taste change on a single good in a nonparametric setting. Our theoretical results were used to develop a quadratic programming procedure to recover the minimal intertemporal (and interpersonal) taste heterogeneity required to rationalise observed choices and we have shown that these conditions are equivalent to solving a "latent virtual price" problem.

A censored quantile approach was used to allow for unobserved heterogeneity, censoring, and endogeneity of expenditure in the application of our approach to the consumption of tobacco in the UK over the period 1980 to 2000 where we might expect large shifts in demand due to taste change. Non-separability between tobacco and alcohol consumption was incorporated using a conditional (quantile) demand analysis. We focussed on a single birth cohort aged between 25 and 35 years old in 1980 and recovered tastes for high and low education individuals at a set of quantiles of the tobacco distribution.

Systematic taste change was required to rationalise tobacco consumption except at the heaviest-smoking quantile of the low-education group, who increased their consumption of tobacco over the period of interest, and at the lightest-smoking quantile of the high-education group, who had effectively ceased smoking by 1986 (the first year that negative taste shifts could be nonparametrically identified in our dataset). A series of negative perturbations to the marginal utility of tobacco were found to be sufficient to rationalise the trends in consumption for all other individuals. For those experiencing negative taste change, we find that the virtual price of tobacco rose over the period 1986-1992 even though the observed relative price declined over the period.

Statistically significant educational differences in the marginal willingness to pay for tobacco were recovered; more highly educated individuals experienced a greater shift in their effective tastes away from tobacco. However, we found that educational differences in tastes narrowed over the period for 'lighter' smokers (.65 and .70 quantiles), while the tastes of high and low education heavy smokers (.85 and .90 quantiles) diverged. We also recovered systematic differences in taste changes of light and heavy drinkers. Focussing on the low-education individuals only for reasons of sample size, we found no evidence of any significant taste changes for heavy consumers of alcohol, while tastes declined over the period for light drinkers.

Governments have a choice of policies that they can use to influence consumption patterns. Our results suggest that information programs designed to discourage smoking were not effective at altering tastes for tobacco amongst low-education, heavy smokers. Increasing the price of tobacco could have had a greater impact on the smoking behaviour of this group given the absence of a 'taste tax' disincentivising their consumption. Our findings also motivate the use of further restrictions on the joint consumption of alcohol and tobacco (e.g. bans on smoking in pubs and bars) and on higher alcohol taxation to reduce demand for tobacco. Such policies would exploit the complementarity between alcohol and tobacco to reduce smoking amongst the heavy-smoking heavy-drinking group given the persistence in their tastes for tobacco.[9] In future work, we hope to explore how the taste shifters we recover using our method relate to observable shifts in Government information campaigns on the health consequences of smoking.

[9] Unfortunately, we are not able directly to test the effectiveness of the 'smoking ban' put in place in the U.K. in 2007 as relative prices for tobacco rose over the period, reducing the ability of our nonparametric method to detect negative taste change.

References

Adda, J. and F. Cornaglia (2010). The effect of bans and taxes on passive smoking. *American Economic Journal: Applied Economics 2*(1), 1–32.

Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International economic review 8*(1), 67–77.

Blundell, R., M. Browning, L. Cherchye, I. Crawford, B. De Rock, and F. Vermeulen (2015). Sharp for sarp: nonparametric bounds on counterfactual demands. *American Economic Journal: Microeconomics 7*(1), 43–60.

Blundell, R., D. Kristensen, and R. Matzkin (2014). Bounding quantile demand functions using revealed preference inequalities. *Journal of Econometrics 179*(2), 112–127.

Blundell, R. and J. L. Powell (2007). Censored regression quantiles with endogenous regressors. *Journal of Econometrics 141*(1), 65–83.

Blundell, R. W., M. Browning, and I. A. Crawford (2003). Nonparametric engel curves and revealed preference. *Econometrica 71*(1), 205–240.

Brown, D. J. and R. L. Matzkin (1998). *Estimation of nonparametric functions in simultaneous equations models, with an application to consumer demand.* Cowles Foundation for Research in Economics at Yale University.

Browning, M. and C. Meghir (1991). The effects of male and female labor supply on commodity demands. *Econometrica: Journal of the Econometric Society*, 925–951.

Chernozhukov, V., I. Fernández-Val, and A. E. Kowalski (2015). Quantile regression with censoring and endogeneity. *Journal of Econometrics 186*(1), 201–221.

Chernozhukov, V. and H. Hong (2002). Three-step censored quantile regression and extramarital affairs. *Journal of the American Statistical Association 97*(459), 872–882.

Dollar, K. M., G. G. Homish, L. T. Kozlowski, and K. E. Leonard (2009). Spousal and alcohol-related predictors of smoking cessation: a longitudinal study in a community sample of married couples. *American Journal of Public Health 99*(2), 231–233.

Dong, Y. and S. Shen (2015). Testing for rank invariance or similarity in program evaluation. *Review of Economics and Statistics* (0).

Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica 17*(66), 159–174.

Imbens, G. W. and W. K. Newey (2009). Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica 77*(5), 1481–1512.

Koenker, R. and G. Bassett Jr (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.

Matzkin, R. L. (2007). Heterogeneous choice. *Econometric Society Monographs 43*, 75.

McFadden, D. L. and M. Fosgerau (2012). A theory of the perturbed consumer with general budgets. Technical report, National Bureau of Economic Research.

Milgrom, P. and C. Shannon (1994). Monotone comparative statics. *Econometrica: Journal of the Econometric Society*, 157–180.

Neary, J. P. and K. W. Roberts (1980). The theory of household behaviour under rationing. *European economic review 13*(1), 25–42.

Rose, H. (1958). Consistency of preference: the two-commodity case. *The Review of Economic Studies 25*(2), 124–125.

Rothbarth, E. (1941). The measurement of changes in real income under conditions of rationing. *The Review of Economic Studies 8*(2), 100–107.

Samuelson, P. A. (1948). Consumption theory in terms of revealed preference. *Economica 15*(60), 243–253.

Varian, H. R. (1982). The nonparametric approach to demand analysis. *Econometrica: Journal of the Econometric Society*, 945–973.

Varian, H. R. (1983). Non-parametric tests of consumer behaviour. *The review of economic studies 50*(1), 99–110.

Varian, H. R. (1988). Revealed preference with a subset of goods. *Journal of Economic Theory 46*(1), 179–185.

Appendix A: Proofs

Theorem 1. The following statements are equivalent:

1. Observed choice behaviour, $D = \{\mathbf{p}_t, \mathbf{q}_t^i\}_{t=1,\ldots,T}$, can be good-1 taste rationalised.

2. One can find sets $\{v_t\}_{t=1,\ldots,T}$, $\{\alpha_t^1\}_{t=1,\ldots,T}$ and $\{\lambda_t\}_{t=1,\ldots,T}$ with $\lambda_t > 0$ for all $t = 1, ..., T$, such that there exists a non-empty solution set to the following revealed preference inequalities:

$$\begin{aligned} v_s - v_t + \alpha_t^1(q_s^1 - q_t^1) &\leq \lambda_t \mathbf{p}_t'(\mathbf{q}_s - \mathbf{q}_t) \\ \alpha_t^1 &\leq \lambda_t p_t^1 \end{aligned} \tag{23}$$

Proof:

Necessity: Let us consider the case where our data set has been generated by the model. Observed choices are then the solution to the following optimisation problem:

$$\max_{\{\mathbf{q}_t\}_{t=1,\ldots,T}} v(\mathbf{q}_t) + \alpha_t^1 q_t^1$$

subject to

$$\mathbf{p}_t' \mathbf{q}_t \leq x_t$$

An optimal solution to the problem must satisfy:

$$\begin{aligned} \nabla_{q_t^1} v(\mathbf{q}_t) + \alpha_t^1 &\leq \lambda_t p_t^1 \\ \nabla_{\mathbf{q}_t^{\neg 1}} v(\mathbf{q}_t) &\leq \lambda_t \mathbf{p}_t^{\neg 1} \end{aligned}$$

Given a particular level of the taste shifter, α_t^1, concavity of the utility function implies:

$$u(\mathbf{q}_t, \alpha_t^1) + \nabla_{\mathbf{q}_t} u(\mathbf{q}_t, \alpha_t^1)'(\mathbf{q}_s - \mathbf{q}_t) \geq u(\mathbf{q}_s, \alpha_t^1)$$

Substituting the first order conditions into the concavity condition and rearranging gives:

$$v(\mathbf{q}_s) - v(\mathbf{q}_t) + \alpha_t^1(q_s^1 - q_t^1) \leq \lambda_t \mathbf{p}_t'(\mathbf{q}_s - \mathbf{q}_t)$$

Letting $v_t = v(\mathbf{q}_t)$, returns the first set of inequalities.

The second set of inequalities are required for the base utility function to be strictly increasing in $\mathbf{q}$. From the first order conditions,

$$\alpha_t^1 \quad > \quad \lambda_t p_t^1$$

would imply

$$\nabla_{q_t^1} v(\mathbf{q}_t) \quad < \quad 0$$

violating monotonicity.

Sufficiency: The concavity condition associated with the taste-varying utility function, $u(\mathbf{q}, \alpha_t^1)$ implies the existence of T overestimates of the utility of some bundle $\mathbf{q}$:

$$\begin{aligned} v(\mathbf{q}) &\leq v_t + \lambda_t \mathbf{p}_t'(\mathbf{q} - \mathbf{q}_t) - \alpha_t^1(q^1 - q_t^1) \\ v(\mathbf{q}) &\leq v_t + \lambda_t \widetilde{\mathbf{p}}_t'(\mathbf{q} - \mathbf{q}_t) \end{aligned}$$

where $\widetilde{p}_t^1 = p_t^1 - \alpha_t^1/\lambda_t$ and $\widetilde{\mathbf{p}}_t^{\neg 1} = \mathbf{p}$. A piecewise linear utility function can be derived from the lower envelope of the hyperplanes defined by these T overestimates:

$$v(\mathbf{q}) = \min_t \{v_t + \lambda_t \widetilde{\mathbf{p}}_t'(\mathbf{q} - \mathbf{q}_t)\}$$

The utility of any feasible consumption bundle cannot be strictly greater than that conferred by observed choices with the utility function defined as above. Consider an arbitrary feasible bundle, $\widehat{\mathbf{q}}$:

$$\mathbf{p}_t'\widehat{\mathbf{q}} \leq \mathbf{p}_t'\mathbf{q}_t$$

Given the definition of the base individual utility function:

$$v(\widehat{\mathbf{q}}) + \alpha_t^1 \widehat{q}^1 \leq v_t + \lambda_t \widetilde{\mathbf{p}}_t'(\widehat{\mathbf{q}} - \mathbf{q}_t) + \alpha_t^1 \widehat{q}_t^1$$

Noting that

$$\lambda_t \widetilde{\mathbf{p}}_t'(\widehat{\mathbf{q}} - \mathbf{q}_t) = \lambda_t \mathbf{p}_t(\widehat{\mathbf{q}} - \mathbf{q}_t) - \alpha_t^1(\widehat{q}^1 - q_t^1)$$

returns

$$\begin{aligned} v(\widehat{\mathbf{q}}) + \alpha_t^1 \widehat{q}^1 &\leq v_t + \lambda_t \mathbf{p}_t'(\widehat{\mathbf{q}} - \mathbf{q}) + \alpha_t^1 q_t^1 \\ &\leq u(\mathbf{q}_t, \alpha_t^1) + \lambda_t \mathbf{p}_t'(\widehat{\mathbf{q}} - \mathbf{q}_t) \end{aligned}$$

Further noting that

$$\begin{aligned} \mathbf{p}_t'\widehat{\mathbf{q}} &\leq \mathbf{p}_t'\mathbf{q}_t \\ \mathbf{p}_t'(\widehat{\mathbf{q}} - \mathbf{q}_t) &\leq 0 \end{aligned}$$

Implies that

$$v(\widehat{\mathbf{q}}) + \alpha_t^1 \widehat{q}^1 \leq v(\mathbf{q}_t) + \alpha_t^1 q_t^1$$

In words, any other feasible bundle yields weakly lower utility than $\mathbf{q}_t$. Therefore, we can always construct a utility function which taste rationalises the data set given that a non-empty solution set is associated with the inequalities of Theorem 1. ■

Theorem 2. Given $q_t^1 \neq q_s^1$ for all t and s, any data set D can be good-1 taste rationalised.

Proof.

The inequalities of Theorem 1 can be expressed in terms of virtual prices.

$$\begin{aligned} v_s - v_t + \alpha_t^1(q_s^1 - q_t^1) &\leq \lambda_t \mathbf{p}_t'(\mathbf{q}_s - \mathbf{q}_t) \\ v_s - v_t &\leq \lambda_t \widetilde{\mathbf{p}}_t'(\mathbf{q}_s - \mathbf{q}_t) \end{aligned}$$

where

$$\begin{aligned} \widetilde{\mathbf{p}}_t &= \begin{bmatrix} p_t^1 - \alpha_t^{1i}/\lambda_t^i \\ \mathbf{p}_t^{\neg 1} \end{bmatrix} \\ \widetilde{p}_t^1 &\geq 0 \end{aligned}$$

Varian (1982) proves that the following conditions are equivalent.

1. A data set $\{\widetilde{\mathbf{p}}_t, \mathbf{q}_t^i\}_{t=1,...,T}$ satisfies GARP.

2. There exist numbers $\{v_t\}_{t=1,...,T}$, $\{\alpha_t^1\}_{t=1,...,T}$ and $\{\lambda_t\}_{t=1,...,T}$ with $\lambda_t > 0$ for all $t = 1, ..., T$ such that the following "Afriat" inequalities hold.

$$v_s - v_t \leq \lambda_t \widetilde{\mathbf{p}}_t^{'}(\mathbf{q}_s - \mathbf{q}_t)$$

We first establish the existence of rationalising shadow prices $\widetilde{\mathbf{p}}_t$. We observe the data set $\{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$. Assume that good-1 exhibits perfect intertemporal variation, i.e. $q_t^1 \neq q_s^1$ for all $t \neq s$. We proceed by extending Theorem 1 from Varian (1988) to the current setting.

If $\{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ satisfies GARP, then the choice set satisfies the inequalities of Theorem 1 with:

$$\alpha_t^1 = \mathbf{0}$$

for $t = 1, ..., T$. If $\{\mathbf{p}_t, \mathbf{q}_t\}_{t=1,...,T}$ fails GARP, then there exist periods s and t such that

$$\begin{aligned} \mathbf{p}_s'\mathbf{q}_s &\leq \mathbf{p}_s'\mathbf{q}_t \\ \mathbf{p}_t'\mathbf{q}_t &\leq \mathbf{p}_t'\mathbf{q}_s \end{aligned}$$

Given perfect intertemporal variation of good-1, there always exists a set of modifications to p_t^1 such that the GARP inequalities are satisfied. This result follows from Theorem 1 in Varian (1988), in which it is proved that, given perfect intertemporal variation for a good with a missing price, one can always find a price trajectory for this good such that the entire data set satisfies GARP. To demonstrate the relevance of Varian (1988) result in this context, let us consider what value p_t^1 would have to take on, if we were to conjecture that, once taste change is taken into account, the consumer prefers the bundle $\mathbf{q}_t$ to $\mathbf{q}_s$. This conjecture implies the need to prove the existence of a price $\widetilde{p}_t^1$ such that:

$$\begin{aligned}
\mathbf{p}_t^{\neg 1\prime}\mathbf{q}_t^{\neg 1} + \widetilde{p}_t^1 q_t^1 &\geq \mathbf{p}_t^{\neg 1\prime}\mathbf{q}_s^{\neg 1} + \widetilde{p}_t^1 q_s^1 \\
\widetilde{p}_t^1 &\geq \frac{\mathbf{p}_t^{\neg 1\prime}(\mathbf{q}_s^{\neg 1} - \mathbf{q}_t^{\neg 1})}{q_t^1 - q_s^1}
\end{aligned}$$

where $\mathbf{p}_t^{\neg 1}$ gives the price vector for all goods except good-1, $\mathbf{p}_t^{\neg 1} = [p_t^2, ..., p_t^K]$, and $\mathbf{q}_t^{\neg 1}$ is defined analogously. For each period t, define the lower bound on the virtual price of good 1 such that:

$$\widetilde{p}_t^1 > \max_{s \neq t} \left\{ \frac{\mathbf{p}_t^{\neg 1\prime}(\mathbf{q}_s^{\neg 1} - \mathbf{q}_t^{\neg 1})}{q_t^1 - q_s^1}, 1 \right\}$$

Further define the "taste adjusted direct revealed preferred relation", $\widetilde{\mathbb{R}}^0$. If $\widetilde{\mathbf{p}}_t^{\prime}\mathbf{q}_t \geq \widetilde{\mathbf{p}}_t^{\prime}\mathbf{q}_s$, then we conclude that $\mathbf{q}_t$ is directly revealed taste preferred to $\mathbf{q}_s$, or $\mathbf{q}_t\widetilde{\mathbb{R}}^0\mathbf{q}_s$. There are then two cases to consider:

1. $q_t^1 > q_s^1$: In this case, we must have that

$$\begin{aligned}
\widetilde{p}_t^1(q_t^1 - q_s^1) &> \mathbf{p}_t^{\neg 1\prime}(\mathbf{q}_s^{\neg 1} - \mathbf{q}_t^{\neg 1}) \\
\widetilde{p}_t^1 q_t^1 + \mathbf{p}_t^{\neg 1\prime}\mathbf{q}_t^{\neg 1} &> \widetilde{p}_t^1 q_s^1 + \mathbf{p}_t^{\neg 1\prime}\mathbf{q}_s^{\neg 1} \\
\widetilde{\mathbf{p}}_t^{\prime}\mathbf{q}_t &> \widetilde{\mathbf{p}}_t^{\prime}\mathbf{q}_s
\end{aligned}$$

   and set $\mathbf{q}_t\widetilde{\mathbb{R}}^0\mathbf{q}_s$.

2. $q_t^1 < q_s^1$. In this case, we must have that

$$\begin{aligned}
\widetilde{p}_t^1(q_t^1 - q_s^1) &< \mathbf{p}_t^{\neg 1\prime}(\mathbf{q}_s^{\neg 1} - \mathbf{q}_t^{\neg 1}) \\
\widetilde{p}_t^1 q_t^1 + \mathbf{p}_t^{\neg 1\prime}\mathbf{q}_t^{\neg 1} &< \widetilde{p}_t^1 q_s^1 + \mathbf{p}_t^{\neg 1\prime}\mathbf{q}_s^{\neg 1} \\
\widetilde{\mathbf{p}}_t^{\prime}\mathbf{q}_t &> \widetilde{\mathbf{p}}_t^{\prime}\mathbf{q}_s
\end{aligned}$$

and thus it is not the case that $\mathbf{q}_t \widetilde{\mathbb{R}}^0 \mathbf{q}_s$.

Therefore, one can determine the preference ordering of consumption bundles solely by reference to the quantity of good-1 consumed and set the taste adjusted price of good 1 to dominate the impact of revealed preference violations in the unadjusted data set. The choice set $\{\widetilde{\mathbf{p}}_t, \mathbf{q}_t\}_{t=1,...,T}$ then passes GARP. Given the equivalence of GARP and a non-empty feasible set to the standard Afriat inequalities, this result then implies that for for any element of the rationalising price set, $\{\widetilde{\mathbf{p}}_t\}_{t=1,...,T}$, there exist numbers $\{v_t\}_{t=1,...,T}$ and $\{\lambda_t\}_{t=1,...,T}$ with $\lambda_t > 0$ such that the following inequalities hold.

$$v_s - v_t \leq \lambda_t \widetilde{\mathbf{p}}_t^{\prime}(\mathbf{q}_s - \mathbf{q}_t)$$

An element of the set of choice rationalising taste shifters associated with $\{\widetilde{\mathbf{p}}_t\}_{t=1,...,T}$ can then be constructed as:

$$\alpha_t^1 = \lambda_t(p_t^1 - \widetilde{p}_t^1)$$

for $t = 1, ..., T$. The fact that the set of rationalising $\widetilde{p}_t^1$ is unbounded above implies that the taste modification to virtual prices, or equivalently, the change in the marginal willingness to pay for good-1, α_t^1/λ_t, is unbounded below.■

Appendix B: Data

This section provides further details on the data used in our analysis. In the first part of our analysis, the nondurable aggregate is the union of the nondurables and alcohol groups below.

The tobacco, nondurable and alcohol good aggregates are constructed from the following subgroups.

1. **Tobacco**: Cigarettes; Other Tobacco.

2. **Nondurables**: Bread; Cereals; Biscuits; Beef; Lamb; Pork; Bacon; Poultry; Fish; Butter; Oils; Cheese; Eggs; Fresh Milk; Processed Milk; Tea; Coffee; Soft Drinks; Sugar; Sweets; Potatoes; Other Vegetables; Fruit; Other Food; Canteen; Other Snacks; Coal; Electric; Gas; Oil; Household consumables; Pet Care; Postage; Telephone; Domestic Services; Chemical Products; Petrol; Rail Fares; Bus Fares; Other Travel; Toys; Books; Entertainment.

3. **Alcohol**: Beer; Wine; Spirits.