

Prices versus Preferences: Taste Change and Tobacco Consumption

Abi Adams (Oxford and IFS), Richard Blundell (UCL and IFS),
Martin Browning (Oxford and IFS), and Ian Crawford (Oxford and IFS)

December 2013, this draft November 2014*

Abstract

A systematic method for incorporating taste variation into a revealed preference framework for heterogeneous consumers is developed. We create a methodology that enables the recovery of the minimal variation in tastes that are required to rationalise the observed choice patterns. This approach is used to examine the extent to which changes in tobacco consumption have been driven by price changes or by taste changes, and whether the significance of these two channels varies across socioeconomic groups. A censored quantile approach is developed to allow for unobserved heterogeneity and censoring of consumption. Statistically significant educational differences in the marginal willingness to pay for tobacco are recovered. More highly educated cohorts are found to have experienced a greater shift in their effective tastes away from tobacco.

1 Introduction

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

*We would like to thank Laurens Cherchye, Dennis Kristensen, Arthur Lewbel and participants at the June 2014 CAM workshop for helpful comments. The authors also gratefully acknowledge financial support from the Economic and Social Research Council and the ERC. All errors are ours.

efficacy of the different modes is important for designing health policies. In this paper we consider the extent to which changes in tobacco consumption are due to price changes and how much is due to preference change. We also consider how tastes evolve across different education strata.

We develop a systematic method for incorporating taste variation into a revealed preference framework for heterogeneous consumers. Our aim is to create a methodology that enables the recovery of the minimal variation in tastes that are required to rationalise the observed choice patterns. We then use this approach to address the question: How much of the fall in tobacco consumption is due to a rise in the relative price of tobacco and how much can be attributed to taste changes? We also consider how tastes evolve across different socio-economic strata, asking the question: Does education matter? The objective is to understand taste change and to inform policy on the balance between information/health campaigns with tax reform.

To disentangle the effects of price and preference change, we use an Afriat (revealed preference) approach that explicitly allows for the fact that we only observe a finite set of data points. The approach follows in the tradition of Afriat (1967), Diewert (1973) and Varian (1982) and builds upon the more recent work of Blundell et al. (2003, 2008), Browning (1989) and Crawford (2010). This approach has the usual features: Reject/accept answers for rationalisation. Tests for rationalisation involve algorithms that yield an answer in a finite number of steps. If no rejection, set identification of objects of interest. If no rejection, rationalisation is possible with smooth, monotone preferences. We derive a surprising result: we show, under mild assumptions on the characteristics of the 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 also show that the rationalisability question can be reinterpreted as a ‘missing price problem’.

The approach is implemented on household consumer expenditure survey data using RP inequality conditions on the conditional quantile demand functions for tobacco. We allow for non-separability with alcohol consumption.

This paper proceeds as follows. Sections 2 and 3 outline our theoretical framework and derive the necessary and sufficient conditions under which observed behaviour and our model of taste change are consistent. Section 4 develops a quadratic programming methodology that can be applied to uncover the minimal amount of interpersonally comparable taste variation that is necessary to rationalise choice behaviour. Section 5 introduces the data used for our empirical investigation of tobacco consumption in the UK and discusses the construction of quantity sequences

for the psuedo cohorts that we draw from the UK Family Expenditure Survey using censored quantile regression methods. Section 6 applies our method to rationalise the changes in tobacco consumption occurring in the U.K. since 1980. Finally, Section 7 concludes our analysis and considers the implications of our findings for government anti-smoking policy moving forward.

2 Theoretical framework

2.1 Taste changes

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 Brown and Matzkin (1998) and McFadden and Fosgerau (2012).

Given any dataset D made up of a sequence of price-quantity observations for a consumer $D = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1, \dots, T}$ it is clear that we can always find a sequence of taste-shifters $\{\boldsymbol{\alpha}_t\}_{t=1, \dots, T}$ which can rationalise the model. Consider the first order conditions

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

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

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

¹The 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.

where α_t/λ_t represents the innovations in willingness-to-pay or ‘*taste-wedge*’ for every good in every period. The behaviour generated by the model

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

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.

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, Theorem 1) shows that this is indeed 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 (i.e. which introduces as many free parameters as there are observations) is clearly extremely permissive - even given its apparently restrictive additive form. Instead, suppose that we have grounds to believe that the most significant taste changes which have taken place have been confined to a single good (denoted good 1) . Then $\alpha_t = [\alpha_t^1, 0, \dots, 0]'$. This restriction yields 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 $j, k \in \{2, \dots, K\}$ is invariant to taste instability on good-1. Another implication of this functional form is that preferences will obey the single crossing property.

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

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

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 for the good of interest - the MRS with respect to the good of interest given α_t^1 is always high (lower) in one period than it is in another at every point in commodity space. Another basic result that follows from single crossing is that $\text{argmax}_{\mathbf{q}} u(\mathbf{q}; \alpha^1)$ increases with α^1 (see Milgrom and Shannon (1994)). We will use this feature of the model to motivate the use of quantile regression methods in the empirical section of this paper.

2.2 Testable conditions

We are interested in establishing whether a consumer's observed choice behaviour, $D = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1, \dots, T}$, 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, $D = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1, \dots, T}$, 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, $\{\alpha_t^1\}_{t=1, \dots, T}$, 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, D can be rationalised by the model if there exists a time-invariant base utility function, $v(\mathbf{q})$ and a series of perturbations to the marginal utility of good-1 such that observed choices are weakly preferred to all feasible alternatives. The testable empirical conditions, involving only observables, that are equivalent to a rationalisation of D by our theoretical model are given in Theorem 1.

Theorem 1. The following statements are equivalent:

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

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

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

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 $\{u_t\}_{t=1,\dots,T}$ and $\{\lambda_t\}_{t=1,\dots,T}$ 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 $t = 1, \dots, T$ returns the standard Afriat inequalities. If there is no intertemporal variation in good-1, $q_t^1 = q_s^1$ for all $t, s \in 1, \dots, T$, then we immediately have the standard Afriat conditions.

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

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

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

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

Note that the shadow price of consumption in period t , given tastes in period s , may be negative if the taste-change term $(\alpha_t^1 - \alpha_s^1)$ 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

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

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 $t - 1$.

The revealed preference inequalities can be given a more intuitive behavioural interpretation when they are expressed in the language of 'virtual budgets'. The concept of a virtual budget was first suggested by Robarath (1940) and Hicks (1941) to develop the theory of choice behaviour under rationing. In our setting, the virtual budget of the consumer at t , $\{\tilde{\mathbf{p}}_t, \tilde{x}\}$, is that which induces the consumer with the utility function $v(\mathbf{q})$ to demand the same bundle as the consumer with the utility function $v^i(\mathbf{q}) + \alpha_t^{1i}$ who faces market prices \mathbf{p}_t and total expenditure x .

The inequalities defined by Theorem 1 can be expressed as:

$$v_s - v_t \leq \lambda_t \tilde{\mathbf{p}}_t'(\mathbf{q}_s - \mathbf{q}_t) \quad (7)$$

$$\tilde{\mathbf{p}}_t \geq 0 \quad (8)$$

where

$$\tilde{\mathbf{p}}_t = \begin{bmatrix} p_t^1 - \alpha_t^1/\lambda_t \\ \mathbf{p}_t^{-1} \end{bmatrix}$$

\tilde{p}_t^1 is to be interpreted as the 'virtual price' of good-1, which is determined by the consumer's tastes at t . The equivalent virtual budget characterisation of the consumer's choice problem is then:

$$\max_{\mathbf{q}} v(\mathbf{q}) \quad s.t. \quad \tilde{\mathbf{p}}_t' \mathbf{q} = \tilde{x} \quad (9)$$

where $\tilde{x} = \mathbf{p}_t' \mathbf{q} - q^1 \alpha_t^1 / \lambda_t$.

Given this 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 changes in 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. 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.

2.3 Testability

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

Definition There is perfect variation in good- k if $q_t^k \neq q_s^k$ for all $t, s = 1, \dots, T$.

Theorem 2 If good-1 satisfies perfect variation in a data set D , then it can be good-1 taste rationalised.

Note that perfect variation is sufficient, but not necessary, for D to be good-1 taste rationalised. For subsets $S_t \subseteq \{1, \dots, T\}$ within which $q_t^1 = q$ for all $t \in S_t$, if the choice set $\{\mathbf{p}_t^{-1}, \mathbf{q}_t^{-1i}\}_{t \in S_t}$ 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 would be if the 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.

2.4 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 the minimal marginal utility perturbations on good-1 that are

necessary 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 $D = \{\mathbf{p}_t, \mathbf{q}_t\}_{t=1, \dots, T}$ are identified as the unique solution set $\{\alpha_t^1\}_{t=1, \dots, T}$ to the following quadratic programme.

$$\min_{\{v_t, \lambda_t, \alpha_t\}_{t=1, \dots, T}} \sum_{t=1}^T (\alpha_t^1)^2 \quad (10)$$

subject to:

1. The revealed preference inequalities:

$$\begin{aligned} v_s - v_t + \alpha_t^1(q_s^1 - q_t^1) &\leq \lambda_t \mathbf{p}'_t(\mathbf{q}_s - \mathbf{q}_t) \\ \alpha_t^1 &\leq \lambda_t p_t^1 \end{aligned} \quad (11)$$

2. The normalisation conditions:

$$\begin{aligned} v_1 &= \delta \quad (\text{an arbitrary constant}) \\ \lambda_1 &= \beta \quad (\text{an arbitrary constant}) \\ \alpha_1^1 &= 0 \end{aligned} \quad (12)$$

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

Minimizing 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 so given ordinality of the utility function. Let $\{\bar{v}_t, \bar{\lambda}_t, \bar{\alpha}_t^1\}_{t=1, \dots, T}$ represent some feasible solution to the rationalisation constraints. As the set represents a feasible solution, the following inequalities hold for all $s, t \in \{1, \dots, T\}$:

$$\bar{v}_s - \bar{v}_t + \bar{\alpha}_t^1(q_s^1 - q_t^1) \geq \bar{\lambda}_t \mathbf{p}'_t(\mathbf{q}_s - \mathbf{q}_t) \quad (13)$$

However, the following set of inequalities is also feasible:

$$\beta(\bar{v}_s + \delta) - \beta(\bar{v}_t + \delta) + \beta \bar{\alpha}_t^1(q_s^1 - q_t^1) \geq \beta \bar{\lambda}_t \mathbf{p}'_t(\mathbf{q}_s - \mathbf{q}_t) \quad (14)$$

for $\delta \in (-\infty, \infty)$ and $\beta > 0$. Thus, without a location normalisation, for example, $v_1 = \delta$, the quadratic programming procedure is ill-posed as 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 $\{\alpha_t^1\}_{t=2,\dots,T}$ as the minimal rationalising marginal utility perturbations to good-1 relative to preferences at $t = 1$. We also impose the scale normalisation $\lambda_1 = \beta$ for computational reasons. This normalisation ensures that the output of the quadratic programming procedure is scaled sensibly and reduces the sensitivity of the solution to the termination tolerance level of the algorithm that is used to solve the programming problem.

3 Interpersonal heterogeneity

In conventional revealed preference analyses, testing and recovery exercises are typically conducted on an individual-by-individual basis. However, following such a strategy is problematic in our context because, in the empirical application conducted in Section 4, we will be interested in comparing the magnitude of taste changes across individuals from different socioeconomic groups. Conducting such an exercise requires us to be able to meaningfully compare the values of α_t^{1i} and α_s^{1j} for any individuals $i, j = \{1, \dots, N\}$ and time periods $s, t = \{1, \dots, T\}$.

It is difficult to compare the magnitude of α^1 parameters gained from independent quadratic programmes because they represent marginal utility perturbations and their magnitude is meaningless independent of some ordinal representation of base preferences. Running our recovery exercise independently for separate individuals does not impose commonality of base preferences and thus the finding that $\alpha_t^{1i} > \alpha_t^{1j}$ does not justify the conclusion that 'individual i 's taste for good-1 was greater than individual j 's taste for good-1'. Further, differences in the base utility function across individuals could lead to an interpersonal violation of single-crossing. Such violations preclude one from making global statements about taste difference for good-1 across individuals.

Extending our theoretical framework to impose commonality of the base utility function provides a solution to both of these problems. Imposing $v^i(\mathbf{q}) = v^j(\mathbf{q}) = v(\mathbf{q})$ for all $i, j \in \{1, \dots, N\}$ ensures a common ordinal representation of base preferences and that interpersonal and intertemporal preferences satisfy the single crossing property in $(\mathbf{q}; \alpha^1)$ space. This facilitates meaningful interpersonal comparison of taste shifters and allows one to make global statements concerning relative tastes for good-1.

Recovery of α^1 (Common base utility) The minimal squared perturbations to the marginal utility of good-1 relative to preferences in period 1 that are necessary to good-1 rationalise the observed choice behaviour $\{D^i\}_{i=1,\dots,N}$,

where $D^i = \{\mathbf{p}_t, \mathbf{q}_t^i\}_{t=1, \dots, T}$, of a set of individuals $i = \{1, \dots, N\}$ when one assumes a common base utility function are identified as the unique solution set $\{\alpha_t^1\}_{t=1, \dots, T}$ to the following quadratic programme.

$$\min_{\{v_t^i, \lambda_t^j, \alpha_t^{i1}\}_{t=1, \dots, T}} \sum_{i=1}^N \sum_{t=1}^T (\alpha_t^{i1})^2 \quad (15)$$

subject to:

1. The revealed preference inequalities:

$$\begin{aligned} v_s^i - v_t^j + \alpha_t^{j1}(q_s^{i1} - q_t^{j1}) &\leq \lambda_t^j \mathbf{p}'_t(\mathbf{q}_s^i - \mathbf{q}_t^j) \\ \alpha_t^{i1} &\leq \lambda_t^i p_t^1 \end{aligned} \quad (16)$$

2. The normalisation conditions:

$$\begin{aligned} v_1^1 &= \delta \\ \lambda_1^1 &= \beta \\ \alpha_1^{1,1} &= 0 \end{aligned} \quad (17)$$

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

In this normalisation everything is relative to agent 1 in period 1. Thus we can meaningfully state that agent $i \neq 1$ initially likes good 1 more than the reference agent if $\alpha_1^{i,1} > 0$. Alternatively they could start out with the same tastes ($\alpha_1^{i,1} = 0$) but have differently evolving tastes ($\alpha_t^{i,1} \neq \alpha_t^{1,1}$ for some $t > 1$). The location normalisation again guarantees that the problem is well-posed and allows us to interpret the recovered α 's as the minimum taste shifts relative to individual 1's tastes at $t = 1$ that are necessary to rationalise observed behaviour.

4 Rationalising tobacco consumption

We now use our revealed preference methodology to recover the minimal taste perturbations for tobacco that are required to rationalise the consumptions patterns of a set of psuedo-cohorts that were stratified by education level and the level of tobacco consumption.


4.1 The data

We draw upon repeated cross-section data from the U.K. Family Expenditure Survey (FES) between 1980 and 2000. The FES records detailed expenditure and demographic information for 7,000 randomly selected households each

year. From this data set we construct two groups of individuals that are stratified by education level and who were aged between 25 and 35 years old in 1980.² The "low" education group, L , is formed from those individuals who left school no older than 16 years old, the legal minimum. Those staying in school past age 16 comprise the "high education" group, H . Appendix A provides further details on the number of observations contributing to each education group over time and additional summary statistics.

We are primarily concerned with choice over a tobacco aggregate and a nondurable commodity aggregate, i.e. $K = 2$. Appendix A 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 displayed in Figure 3, the relative price of tobacco relative to nondurables rose significantly between 1980 and 2000. The rise was smooth and systematic, apart from a brief stagnation in the late 1980s that was caused by a reduction in growth of the absolute price of tobacco rather than an increase in the growth rate of nondurable prices.

Figure 3. Relative price of tobacco

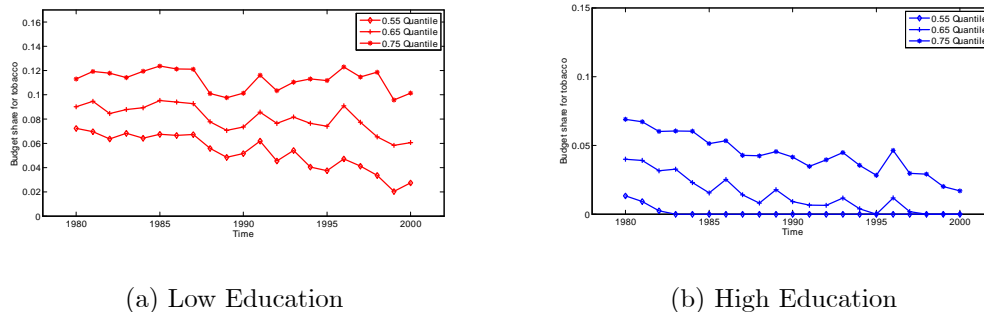


In the following sections, we will rationalise the changes in tobacco consumption at three different quantiles of the budget share for tobacco distribution, $\varepsilon = \{0.55, 0.65, 0.75\}$, for each education group. We refer to the 0.55-quantile as "light smokers", to the 0.65-quantile as "moderate smokers" and we refer to the 0.75-quantile as "heavy smokers". We thus recover taste changes for six different cohorts. Figure 4 displays the budget shares for tobacco at the relevant quantiles of both education cohorts between 1980 and 2000. At each quantile, the high education cohort devotes a smaller proportion of their budget to tobacco than the low education cohort. In fact, within the high education

²We select this birth cohort because less than 5% of smokers start smoking after they reach their 25th birthday (Office for National Statistics, 2012). The assumption that the population of smokers is stable is then relatively mild.

cohort, the 0.55 quantile of the budget share for tobacco distribution falls to zero in 1983 and remains nil for the rest of the period considered, while the 0.65 quantile hits zero by the end of the period. Tobacco consumption remains strictly positive in all periods for the quantiles of the low education cohort that are considered.

Figure 4. Budget share for Tobacco



4.2 Sequential maximum power paths

We are interested in recovering the contribution of taste change to the trends in tobacco consumption over time and between different cohorts. However, budget variation in the data hinders our progress somewhat. Income growth in survey data often prevents the intersection of budget hyperplanes, which weakens the ability of our method to detect violations of stable preferences. Thus, to make our methodology as powerful as possible, we evaluate each cohort's demands along the Blundell, Browning and Crawford (2003) "Sequential Maximum Power" (SMP) paths. Intuitively, the SMP method controls for budget variation to provide us with a choice set that maximises the chance of detecting violations of a time-invariant utility function in observational data.

To illustrate the key ideas lying behind the SMP procedure, consider Figure (the SMP picture from the BBC 2003 & 2008 papers, TBC)

Figure 5: SMP path

In our empirical application, we recover the minimal α perturbations that are required to rationalise choice behaviour along the SMP path that starts at an education group's median demand in 1980, and which then continues sequentially over time to select the demand that is just weakly preferred to the SMP demand in the previous period. Specifically, we recover taste changes for cohort demands at the SMP expenditure levels $\{\tilde{x}_t\}_{t=1,\dots,T}$, which are determined as:

$$x_1^c = Q_x(0.5|c) \tag{18}$$

and

$$\tilde{x}_t^c = \mathbf{p}_t' \mathbf{q}_{t-1}^c \quad (19)$$

for $t = 2, \dots, T$ and $c = \{L, H\}$, where $\mathbf{q}_t^c = \mathbf{q}^c(\mathbf{p}_t, \tilde{x}_t)$.³

We can abstract from the complicating issues caused by transitivity in the construction of the SMP path that are examined in Blundell *et al.* (2012) because ours is a two-good demand system. As first highlighted by Rose (1958), transitivity has no empirical content for a 2-good demand system.

4.3 Censored quantile regression

To operationalise the SMP methodology, we require the sequence of quantities $\{\mathbf{q}^c(\mathbf{p}_t, \tilde{x}_t)\}_{t=1, \dots, T}$. However, we only observe $\mathbf{q}^c(\mathbf{p}_t, x_t)$ directly. The implementation of the SMP framework thus requires us to estimate the expansion paths for each education cohort at each price regime. To do so we employ censored quantile regression methods in the spirit of Chernozhukov and Hong (2002) and Chernozhukov, Fernandez-Val and Kowalski (2010). Censored quantile regression models are a natural choice given our theoretical framework; with single crossing, the quantile rank of an individual's α parameter identifies the quantile rank of their tobacco demand. Further, the latent budget share for tobacco, w^{i*} is left censored at zero: $w^i = \max\{0, w^{i*}\}$. Quantile regression models are equivariant to monotone transformations such as censoring and thus allow us to avoid the strong parametric assumptions of traditional Tobit estimators.

Formally, we consider the following triangular system of quantile equations for the budget share of tobacco, w^i , at each price regime:

$$\begin{aligned} w^i &= \max(0, w^{i,*}) \\ w^{i,*} &= Q_{w^{i,*}}(\epsilon^i | x^i, \mathbf{z}^i, v^i, c^i) \\ x^i &= Q_{x^i}(v^i | z^i, m^i, c^i) \end{aligned} \quad (20)$$

where

$$\begin{aligned} \epsilon^i &\sim U(0, 1) | x^i, \mathbf{z}^i, v^i, m^i, c^i \\ v^i &\sim U(0, 1) | \mathbf{z}^i, m^i, c^i \end{aligned} \quad (21)$$

³Please note that demands for light, moderate and heavy smokers within each education group are recovered at the same expenditure level, but that the SMP expenditure level differs between education groups.

and x^i is total expenditure (on nondurables and tobacco), \mathbf{z}^i is a vector of household characteristics, v^i is an unobserved latent variable that is included to account for the possible endogeneity of x^i , and $c^i \in \{L, H\}$ denotes individual i 's education cohort membership. m^i , the log of disposable income, is our excluded instrument that allows us to recover v^i as the conditional quantile of x given (\mathbf{z}^i, m^i, c^i) .⁴

We use this model to identify conditional cohort demands at the set of quantiles of the tobacco distribution $\epsilon = \{0.55, 0.65, 0.75\}$, for each price regime and education group, along an education group's SMP expenditure path. We note that the estimated model imposes a rank invariance restriction on the budget share for tobacco distribution within education groups; taste evolution and price effects cannot alter the ranking of individuals within the same education group in the budget share for tobacco distribution. Therefore, for any uncensored quantiles of the budget share distribution within a particular education group, ϵ and ϵ' , with $\epsilon > \epsilon'$, we assume that $\alpha_t^\epsilon > \alpha_t^{\epsilon'}$ for all $t = 1, \dots, T$. This assumption is strong but, in our opinion, not incredible. We only consider individuals who are old enough to have made an initial decision to smoke and high profile anti-smoking campaigns often are not targeted in a way that would cause one to expect significant re-ranking within education cohorts.

Estimation To estimate this model, we extend the estimator of Chernozhukov, Fernandez-Val and Kowalski (2010) to allow for unrestricted heterogeneity between education groups.

$$\begin{aligned} Q_{w^{i,*}}(\epsilon^i | x^i, \mathbf{z}^i, v^i, c^i) &= \mathbf{X}^{i'} \boldsymbol{\beta}^{c^i}(\epsilon^i) \\ \mathbf{X}^i &= f(x^i, \mathbf{z}^i, v^i) \\ &= [1, \log(x^i), \log(x^i)^2, oecd^i, v^i] \end{aligned} \tag{22}$$

where $c^i \in C$ and $oecd^i$ is the OECD equivalence index for household i .

Assuming a random sample of N observations at each price regime $\boldsymbol{\beta}^c(\epsilon)$ for $c = \{L, H\}$ is estimated as:

$$\boldsymbol{\beta}^c(\epsilon) = \arg \min_{\boldsymbol{\beta}} \sum_{i \in S^c} \rho_\epsilon(w^i - \hat{\mathbf{X}}^{i'} \boldsymbol{\beta}) \tag{23}$$

where ρ_ϵ is the standard asymmetric absolute loss function of Koenker and Bassett (1978). S^c denotes the set of observations on individuals i for which $c^i = c$ and $Pr(w^i > 0) > \delta$, i.e. the subset of individuals in education group c for which the probability of censoring is negligible and a linear functional form for the conditional quantile is justified.

\mathbf{X}^i is replaced with $\hat{\mathbf{X}}^i = f(x^i, \mathbf{z}^i, \hat{v}^i)$, where \hat{v}^i is an estimator of v^i , because the true value of v^i is unobserved.

⁴See Blundell and Powell (2007) for a similar application.

We assume an additive linear model such that v is easily recovered from the cumulative distribution function of the least squares residuals:

$$\begin{aligned} Q_x(\nu|z, m, c) &= \mathbf{Z}'\boldsymbol{\pi}^c + Q_v(\nu|c) \\ \widehat{v}_i &= Q_v^{-1}(x_i - \mathbf{Z}'_i\widehat{\boldsymbol{\pi}}^c|c) \end{aligned} \tag{24}$$

We refer the reader to Appendix A and Chernozhukov, Fernandez-Val and Kowalski (2010) for a more detailed exposition of the methodology, including the construction of S^c .

4.4 Quadratic programming procedure

We use the quadratic programming procedure that was outlined in Section 2 to recover the minimal set of taste shifters that are required to rationalise the choice behaviour of light, moderate and heavy smokers in each education group at the SMP expenditure levels. Specifically, we recover minimal perturbations to the marginal utility of tobacco for our set of cohorts using the quadratic programme:

$$\min_{\{v_t^{i,\tau}, \lambda_t^{i,\tau}, \alpha_t^{i,\tau}\}} \sum_{\tau \in \boldsymbol{\varepsilon}} \sum_{i \in \{L, H\}} \sum_{t=1}^T \alpha_t^{i,\tau'} \alpha_t^{i,\tau} \tag{25}$$

subject to:

1. The revealed preference inequalities:

$$\begin{aligned} v_s^{i,\tau} - v_t^{j,\delta} + \alpha_t^{j,\delta}(q_s^{i,\tau} - q_t^{j,\delta}) &\leq \lambda_t^{j,\delta} \mathbf{p}'_t(\mathbf{q}_s^{i,\tau} - \mathbf{q}_t^{j,\delta}) \\ \alpha_t^{i,\tau} &\leq \lambda_t^{i,\tau} p_t^1 \end{aligned} \tag{26}$$

for $s, t = \{1980, 1982, 1984, 1986, 1988, 1990, 1992, 1994, 1996, 1998, 2000\}$ and $\tau, \delta \in \{0.55, 0.65, 0.75\}$.

2. The normalisation conditions:

$$\begin{aligned} v_1^{L,0.55} &= 0 \\ \lambda_1^{L,0.55} &= 20 \\ \alpha_1^{L,0.55} &= 0 \end{aligned} \tag{27}$$

Given the normalisation conditions, estimated taste shifters are to be interpreted as perturbations in the marginal utility for tobacco relative to the tastes of the heavy smoking, low education cohort in 1980.

Implementing the quadratic programme above is a computationally intensive procedure. To limit the computational burden of the empirical exercise, we restrict observations to have biennial periodicity. Therefore, each cohort's marginal utility perturbations of tobacco are minimized subject to 4356 rationalisation constraints.

We bootstrap the procedure to address the issue of sampling variation in estimated quantity sequences. To do so, we randomly draw observations with replacement within each education-time cell 1000 times and estimate quantile demands on each resampled set of observations. Minimal taste change is estimated for these sets of perturbed quantities, allowing us to construct a simultaneous 95% confidence interval on taste perturbations for each cohort.

4.5 Recovering virtual prices

As discussed in Section 1.3, the magnitude and pattern of taste changes is more easily interpreted once we cast our framework in the language of virtual budgets. Using the output of the quadratic programme above, we can recover the minimal virtual prices faced by each cohort as:

$$\widehat{\pi}_t^{c,\epsilon} = p_t^1 - \frac{\widehat{\alpha}_t^{c,\epsilon}}{\widehat{\lambda}_t^{c,\epsilon}}$$

The minimum change in marginal willingness to pay, which we term the “taste wedge” is then captured by the term:

$$\widehat{taste_wedge}_t^{c,\epsilon} = \frac{\widehat{\alpha}_t^{c,\epsilon}}{\widehat{\lambda}_t^{c,\epsilon}}$$

5 Results

We find significant differences in the virtual price trajectories along the SMP total expenditure path across cohorts. Figures 6 and 7 depict the minimum virtual prices and the minimal taste wedge that are necessary to rationalise the choice behaviour of the low and high education cohorts, given the normalisation of taste shifters relative to the heavy-smoking low-education cohort in 1980. In addition to recovered virtual prices, the observed relative price for tobacco is depicted. This is the trajectory that each cohort’s virtual price would follow in the absence of intercohort and intertemporal taste heterogeneity. This represents the change in the marginal willingness to pay for tobacco relative to base tastes along the SMP expenditure paths.

Some degree of taste variation is necessary to rationalise the behaviour of every cohort for the period 1980-2000; all virtual price trajectories are significantly different from the price that is observed in reality. Figures 6 and 7 display the intracohort heterogeneity in virtual prices and the taste wedge. We here suppress the 95% confidence intervals to allow for an uncluttered overview of intracohort trends but they are given in later Figures. Unsurprisingly, the virtual price of heavier smoking cohorts is lower compared to that of lighter smoking cohorts. The marginal willingness to

pay for tobacco, captured by the taste wedge, is, therefore, lower for lighter smoking cohorts compared to that dictated by the tastes of the low-education heavy-smokers in 1980.

Figure 6. Virtual price for tobacco

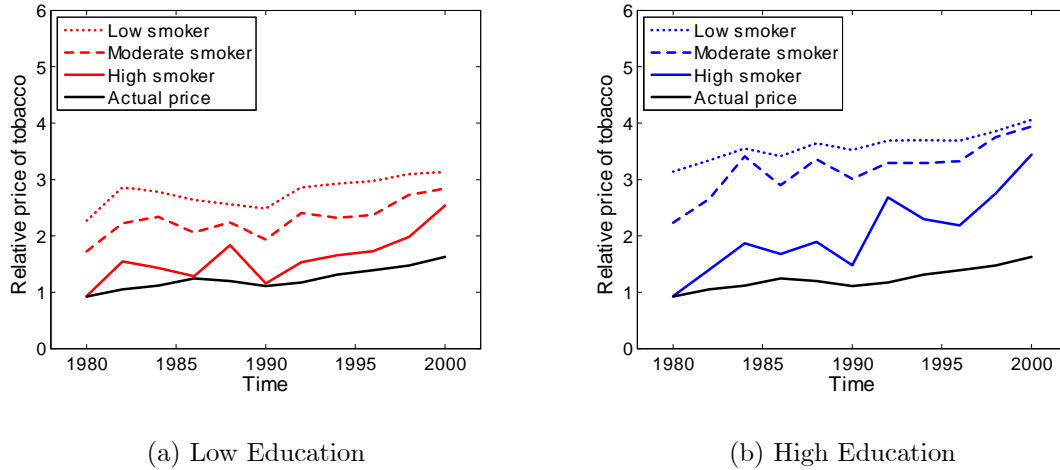
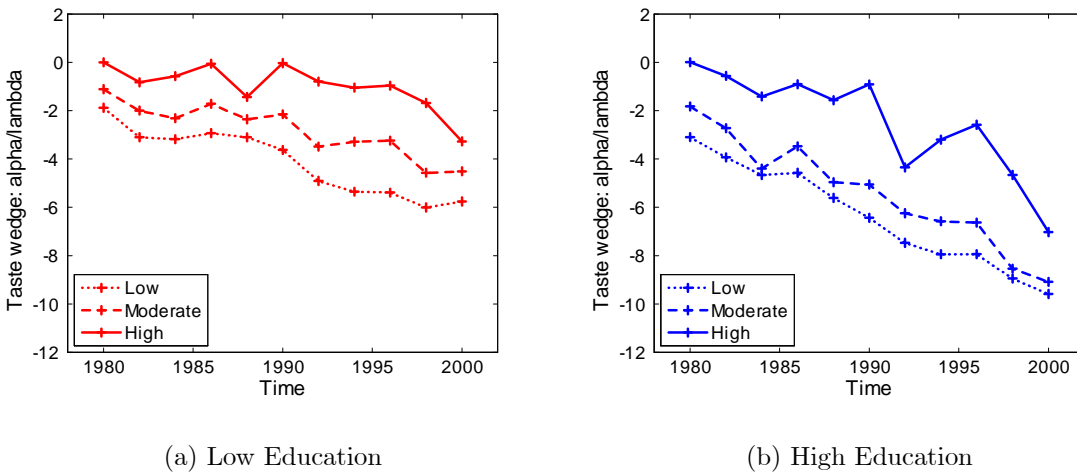


Figure 7. Taste wedge for tobacco




Turning now to compare the virtual price trajectories across education groups, Figures 8 and 9 show the evolution of virtual prices and the taste wedge by smoking group. We find statistically significant heterogeneity in virtual prices across education cohorts for light and moderate smokers along their respective quantity sequences. This is highlighted by reference to Figures 8 and 7, which illustrate that the confidence intervals on the virtual price and taste wedge trajectories for low and high education cohorts are disjoint at light and moderate smoking quantiles.


However, educational differences in virtual prices are not significant for heavy smokers; the 95% confidence intervals over the minimal virtual prices trajectories overlap for the low and high education cohorts at the 0.75-quantile. It should also be noted that there are a number of periods in which the taste wedge of the low-education

heavy-smoking group is not significantly different from zero. These findings suggest that the taste for tobacco is more robust amongst heavier smokers, which accords with biological evidence that nicotine addiction has a genetic basis. Particular gene sequences (especially variants in chromosome 15) have been found to be associated with, among other factors, the number of cigarettes smoked per day and nicotine intake (see, for example, Mansfelder (2000), Berrettini (2008), Keskitalo (2009)). Given that one's genome does not change except by random mutation, one would expect tastes grounded in genetic factors to be relatively time invariant as they appear to be in the case of our heavy-smokers.

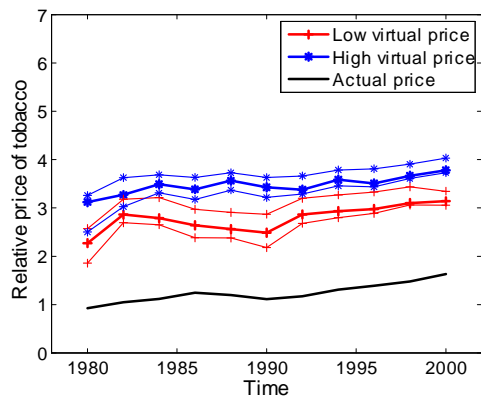
Figure 8. Virtual Price for Tobacco



(a) Heavy smokers




(b) Moderate smokers




(c) Light smokers


Figure 9. Taste wedge



(a) Heavy smokers



(b) Moderate smokers



(c) Light smokers

5.1 Initial conditions and taste evolution


It is interesting to consider the extent to which the intercohort variation in taste wedge that we observe is driven by differences in a cohort’s tastes in 1980 (their ‘initial condition’) and how much due to differences in how tastes have evolved since 1980. To do so we decompose our recovered parameters as follows:

$$\alpha_t^i = \alpha_{1980}^i + \tilde{\alpha}_t^i \quad (28)$$

for $t > 1980$. Note that all recovered taste parameters, α , are normalised relative to the taste of the heaviest smoking, low education cohort in 1980.


Figure 10 gives initial tastes, α_{1980} , for our cohorts. Heterogeneity in initial tastes could be driven by, for example, differences in the extent to which past anti-tobacco campaigns influenced preferences and any genetic variation in tobacco tastes across the tobacco distribution. We see that, unsurprisingly, lighter smoking cohorts have large initial taste differences relative to the low education, heavy smoking cohort and that within the light smoking group, the more highly educated cohort has a stronger dislike for tobacco than the less educated cohort. However, we don’t observe stark differences by education group in the recovered initial tastes for tobacco among moderate and heavy smokers.

Figure 10: Initial taste differences




Turning to consider the evolution of α over time, Figure 11 depicts $\tilde{\alpha}_t$ by smoking group. We observe that, among heavy and moderate smokers, our high educated cohorts typically face a large negative influence to their taste for tobacco. For moderate smokers, there is only one year in which $\alpha_t^{low} < \alpha_t^{high}$. For light smokers, although taste changes are always negative compared to their initial tastes but there is no discernible pattern over time or a clear differences in the evolution of $\tilde{\alpha}_t$ across education cohorts.


Figure 11. $\tilde{\alpha}$ evolution



(a) Heavy smokers



(b) Moderate smokers



(c) Light smokers

6 Conditional taste change

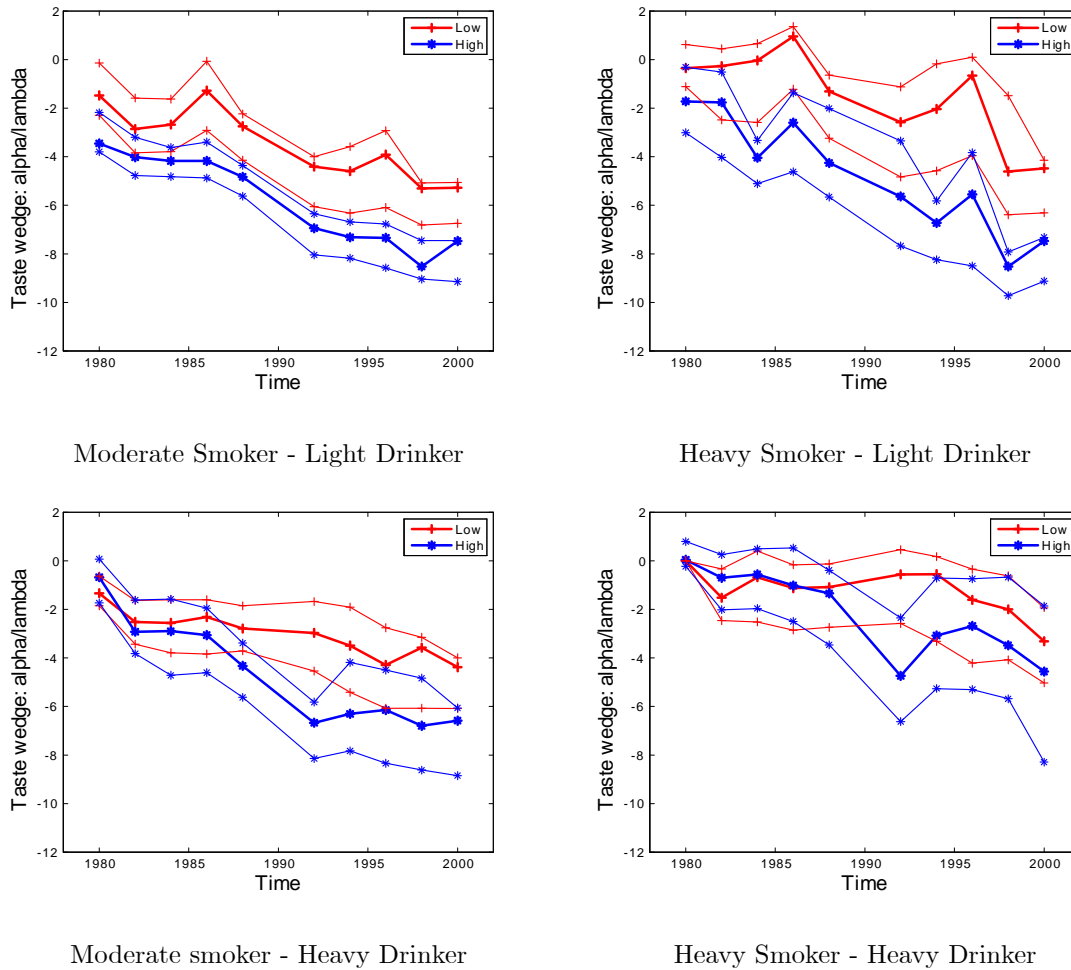
There are a number of separability assumptions that are implicit in our estimation strategy. Not including alcohol as a commodity in our estimated demand system amounts to an assumption concerning the separability of alcohol and tobacco in the utility function. To examine the robustness of our findings on to this assumption, and to explore whether additional patterns emerge from the data once this condition is relaxed, we re-run our quadratic programming procedure on quantile demands that are estimated conditional on alcohol consumption. Specifically, we partition the set of observations that comprise each education group into "light" and "heavy" drinkers depending on whether an individual demands below or above the median budget share for alcohol in that time period. The quantile regression method outlined in the previous section is then applied to estimate demands within each education-alcohol-time cell at the expenditures along the same SMP paths that were calculated previously. Base preferences are normalised relative to those of the heavy smoking, heavy drinking, low education cohort in 1980.

Taste shifters are only estimated for moderate and heavy smokers that are drawn from our education-alcohol cohorts. Estimating comparable taste shift parameters for all quantile, education and alcohol groups yields a quadratic programme with over 17,000 constraints. This is an exceptionally computationally demanding procedure to implement. In addition, the requirement of perfect intertemporal variability of good-1 is violated for demands on the "high education"- "light smoking"- "light drinking" cohort; for most of the period considered the estimated budget share on tobacco for this cohort is zero. We thus do not calculate taste changes for light smokers. This ensures that there will always exist a set of comparable, choice-rationalising taste shifters that can be estimated in a reasonable amount of time. However, please note that the fact that light smokers 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 recognise that the rank invariance assumption, which is imposed by the quantile regression model that is used to estimate SMP demands, is strong in this context. It amounts to a no re-ranking requirement on the joint distribution of tobacco-alcohol group budget shares. We cannot determine how strong this assumption is because we do not have access to panel data for the period. However, research suggests a robust, if modest, positive correlation between alcohol and smoking consumption that is replicated across many studies, which lends some support to our strategy (Bobo and Husten, 2000; Falk, Yi and Hiller-Sturmhofel, 2006).

Results The main themes arising from our earlier results are robust to conditioning of demands upon alcohol consumption. We find our low education cohorts have typically experienced less taste change than high education cohorts in each smoking level-drinking level cell. However, it is only for the moderate smoking- low drinking cohort that educational differences in tastes are statistically significant; in this case, the 95% confidence intervals on virtual prices and the taste wedge are disjoint from 1984 onwards. The marginal willingness to pay evolved very little for the heavy-smoking heavy-drinking cohorts, regardless of education level, as evidenced by the fact that their virtual price trajectories closely follow the observed relative price path. This finding is consistent with government and health practitioner reports that note low smoking cessation rates amongst heavy drinkers (Dollar *et al.*, 2009).

Figure 12: Conditional demands- Taste wedge



7 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. The approach we have developed establishes that we can almost always rationalise a data set when allowing for taste change in a single good. 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 and censoring of consumption. Non-separability between tobacco and alcohol consumption was incorporated using a conditional (quantile) demand analysis. We found that taste change was required for all groups in our expenditure survey data. A series of strictly negative perturbations to the marginal utility of tobacco were found to be sufficient to rationalise the trends in tobacco consumption.

Statistically significant educational differences in the marginal willingness to pay for tobacco were recovered; more highly educated cohorts experienced a greater shift in their effective tastes away from tobacco. We find virtual prices and the taste wedge are disjoint across education groups for all cohorts except for the "heavy smoking"- "heavy drinking" group. Education is irrelevant for explaining the evolution of virtual prices amongst heavy smokers. This might suggest diminished differences in smoking behaviour by education group in the future.

References

- [1] Sydney Afriat (1967), "The Construction of Utility Functions from Expenditure Data", *International Economic Review*, 8(1), 67-77.
- [2] Sydney Afriat (1977), *The Price Index*, London: Cambridge University Press.
- [3] Laura Blow and Ian Crawford (1999), "Valuing Quality", *IFS Working Paper Series W99/21*, The Institute for Fiscal Studies.

- [4] Richard Blundell, Martin Browning and Ian Crawford (2003), "Nonparametric Engel Curves and Revealed Preference", *Econometrica*, 71(1), 205-240.
- [5] Richard Blundell, Martin Browning and Ian Crawford (2008), "Best Nonparametric Bounds on Demand Responses", *Econometrica*, 76(6), 1227-1262.
- [6] Richard Blundell, Martin Browning, Laurens Cherchye, Ian Crawford, Bram De Rock and Frederic Vermeulen (2012), "Sharp for SARP: Nonparametric Bounds on the Behavioural and Welfare Effects of Price Changes", *IFS Working Paper W12/14*, The Institute for Fiscal Studies.
- [7] Richard Blundell and James Powell (2007), "Censored Regression Quantiles with Endogenous Regressors", *Journal of Econometrics*, 141(1), 65-83.
- [8] Janet Bobo and Corinne Husten (2000), "Sociocultural Influences on Smoking and Drinking", *Alcohol Research and Health*, 24(4), 225-232.
- [9] Donald Brown and Rosa Matzkin (1998), "Estimation of Nonparametric Functions in Simultaneous Equations Models, with an Application to Consumer Demand", *Cowles Foundation Discussion Papers 1175*, Cowles Foundation, Yale University.
- [10] Martin Browning (1989), "A Nonparametric Test of the Life-Cycle Rational Expectations Hypothesis", *International Economic Review*, 30(4) 979-992.
- [11] Victor Chernozhukov and Han Hong (2002), "Three-Step Censored Quantile Regression and Extramarital Affairs", *Journal of the American Statistical Association*, 40(459), 872-882.
- [12] Victor Chernozhukov, Iván Fernández-Val and Amanda Kowalski (2010), "Quantile Regression with Censoring and Endogeneity", *Mimeo*, Massachusetts Institute for Technology.
- [13] Ian Crawford (2010), "Habits Revealed", *Review of Economic Studies*, 77(4), 1382-1402.
- [14] Department of Health (2011), *White Paper: Healthy Lives, Healthy People: Our Strategy for Public Health in England*, London: TSO.

- [15] W. Erwin Diewert (1973), "Afriat and Revealed Preference Theory", *The Review of Economic Studies*, 40(3), 419-425.
- [16] Charles Manski (2007), *Identification for Prediction and Decision*, Cambridge, MA: Harvard University Press.
- [17] Daniel McFadden and Mogens Fosgerau (2012), "A Theory of the Perturbed Consumer with General Budgets", *NBER Working Papers 17953*, National Bureau of Economic Research, Inc.
- [18] Paul Milgrom and Chris Shannon (1994), "Monotone Comparative Statics", *Econometrica*, 62(1), 157-180.
- [19] Christopher Murray and Alan Lopez (1997), "Global Mortality, Disability, and the Contribution of Risk Factors: Global Burden of Disease Study", *The Lancet*, 349(9063), 1436-1442.
- [20] Peter Neary and Kevin Roberts (1980), "The Theory of Household Behaviour Under Rationing", *European Economic Review*, 13(1), 25-42.
- [21] James Powell (1984), "Least Absolute Deviations Estimation for the Censored Regression Model", *Journal of Econometrics*, 25(3), 303-325.
- [22] Hugh Rose (1958), "Consistency of Preference: The Two-Commodity Case", *Review of Economic Studies*, 25(2), 124-125.
- [23] Hal Varian (1982), "The Nonparametric Approach to Demand Analysis", *Econometrica*, 50(4), 945-973.
- [24] Hal Varian (1983), "Non-Parametric Tests of Consumer Behaviour", *Review of Economic Studies*, 50(1), 99-110.
- [25] Hal Varian (1988), "Revealed Preference with a Subset of Goods", *Journal of Economic Theory*, 46(1), 179-185.
- [26] World Health Organisation (2012), "Mortality Attributable to Tobacco", *WHO Global Report Series*.