

Income effects and the welfare consequences of tax in differentiated product oligopoly

Rachel Griffith, Lars Nesheim and Martin O’Connell*

December 19, 2016

Abstract

Random utility models are widely used to study consumer choice. The vast majority of applications assume utility is linear in consumption of the outside good, which imposes that total expenditure on the subset of goods of interest does not affect demand for inside goods and restricts demand curvature and pass-through. We show that relaxing these restrictions can be important, particularly if one is interested in the distributional effects of a policy change, even in a market for a small budget share product category. We consider the use of tax policy to lower fat consumption and show that a specific (per unit) tax results in larger reductions than an ad valorem tax, but at greater cost to consumers.

Keywords: income effects, compensating variation, demand estimation, oligopoly, pass-through, fat tax

JEL classification: L13, H20

Acknowledgments: The authors gratefully acknowledge financial support from the European Research Council under ERC-2009-AdG grant agreement number 249529 and from the Economic and Social Research Council under the Centre for the Microeconomic Analysis of Public Policy (grant number RES-544-28-0001), the Centre for Microdata Methods and Practice (grant number RES-589-28-0001), and under the Open Research Area grant number ES/I012222/1 and from ANR under Open Research Area grant number ANR-10-ORAR-009-01.

* Correspondence: Griffith: Institute for Fiscal Studies and University of Manchester, rgriffith@ifs.org.uk; Nesheim: CeMMAP and University College London, l.nesheim@ucl.ac.uk; O’Connell: Institute for Fiscal Studies, martin_o@ifs.org.uk

1 Introduction

Random utility models are widely used to study consumer choice among differentiated products. It is common when using such models to assume that utility is linear in consumption of the outside good. This assumption increases model tractability, simplifies analysis of counterfactual equilibria and simplifies welfare calculations. However, it also places strong restrictions on choice behaviour. For instance, it imposes independence between the total resources a consumer allocates to a set of separable goods and demand for those goods (i.e. it imposes no “income effects”). It also places strong restrictions on the curvature of individual and market demand, and hence on predictions of pass-through of cost shocks and taxes to consumer prices (see Weyl and Fabinger (2013) and Fabinger and Weyl (2015)). Nevertheless, it is commonly believed that for small budget share product categories the restriction that the marginal utility of expenditure is constant is a reasonable approximation.

Our first contribution in this paper is to show that flexibly modelling the impact of total expenditure on demand in a discrete choice demand model can be important, even when interest focuses on a product category that represents a small share of consumer budgets. We allow for a highly flexible relationship between demand and total expenditure and, to highlight the implications, we use our model to simulate the introduction of a tax and compare the implications for demand, tax pass-through and welfare with those implied by specifications that are standard in the literature. We show that this added flexibility is crucial for recovering how the impact of policy reform varies across consumers at different points of the total spending distribution.

Our second contribution is to compare the efficacy of a specific (per unit of saturated fat) and ad valorem tax in reducing consumption of saturated fat. We simulate each tax and account for the possibility that consumers respond by switching across brands, across pack sizes and out of the market. We model the equilibrium pricing response of the multi-product firms in the market assuming that manufacturers set prices in a Nash-Bertrand game. Crucial to the empirical robustness of our estimates is that we relax curvature restrictions that are commonly placed on demand curves and that would otherwise constrain tax pass-through. Our pass-through results are not driven by functional form assumptions.

In our application we focus on the UK butter and margarine market. Our motivation for this is twofold. Firstly, butter and margarine, which makes up a small share

of consumers' grocery spending, is a market where *a priori* we may think abstracting from income effects by assuming a constant marginal utility of expenditure is not very restrictive. Secondly, products in this market are a major source of saturated fat, making the market a leading contender to be subject to a "fat tax". Our model assumes grocery demand is weakly separable from demand for other goods. We consider the decision the consumer makes over which, if any, butter or margarine product to purchase, and how much money to spend on other groceries (the outside good). We model spending on the outside good as entering utility as a flexible polynomial. This allows for the possibility that as a consumer's grocery budget rises, they may alter their butter and margarine demand due to changing marginal utility of grocery expenditure. It also relaxes curvature restrictions commonly placed on demand curves in discrete choice models.

We estimate the model using longitudinal data on the grocery purchases of a representative panel of households. We measure the grocery budget of each household as their mean weekly grocery spending over the year. We control for the structure of households and their labour supply, and use cross-sectional variation in grocery budgets and butter and margarine demands conditional on these demographics. This enables us to isolate the income effect capturing how households will reallocate their grocery budgets between different butter and margarine products and other groceries when their grocery budgets rise. This can be combined with estimates of how total income is allocated to groceries and other broad commodities (see, for instance, Banks et al. (1997)) to get a picture of the impact of income shocks on butter and margarine demands.

We find that the marginal utility of expenditure is not constant – falling convexly in the household's grocery budget. Therefore, as a household's grocery budget rises it will elect to purchase butter and margarine more frequently and to switch to higher quality (and higher priced) varieties. The standard assumption of linear utility is easily rejected. The nonlinearity in utility drives substantial nonlinearity in how consumers are affected by the introduction of a tax. When compared with high expenditure consumers, low expenditure consumers are more responsive to price changes (due to a higher marginal utility of grocery expenditure) and lose out less in monetary terms as a consequence of the introduction of a tax. Therefore low expenditure households respond more strongly to the tax (in percentage terms) than high expenditure house-

holds but the burden they face is less (in monetary terms). Allowing for flexibility in the relationship between grocery budgets and demands is crucial for uncovering these distributional effects of policy reform.

Our second main finding is that, while a specific tax is more effective at lowering saturated fat purchases than an ad valorem tax, it leads to larger losses in consumer welfare. As a consequence the average consumer welfare cost per unit reduction in saturated fat purchases is broadly similar between the two taxes. Although the specific tax more directly targets saturated fat content than the ad valorem tax, it has higher pass-through to consumer prices, leading to higher reductions in consumer welfare.

The existing literature that uses logit models to estimate consumer demand for differentiated goods in grocery markets has typically assumed that utility is linear in total expenditure (or income) (see, Nevo (2001), Villas-Boas (2007)). Under this assumption, expenditure drops out of the model when comparisons are made across alternatives and income effects are therefore ruled out. To capture the cross-sectional relationship between consumers' grocery budgets and purchase patterns, researchers often include a measure of the grocery budget (or income) in a reduced form way as a preference shifter, which linearly shifts the coefficient on price. We show, in our context, that this specification yields estimates of market level average quantities, such as tax pass-through, average price elasticities and aggregate welfare effects, that are similar to those generated by our model. However, it fails to fully recover variation in price sensitivity and welfare effects across the expenditure distribution.

The standard model rules out a causal relationship between a consumer's grocery budget and their butter and margarine demand. However, it does allow for some cross-sectional correlation in demand patterns and welfare effects with total expenditure. We show that when grocery budgets do have a causal impact on demands, but the counterfactual of interest involves relatively small price changes that do not themselves induce large income effects, a simple modification to the standard model, which involves interacting price with higher order expenditure terms, can do a very good job of replicating the distributional results found with our model. The reason for this is that, even though a consumer's utility function may be highly nonlinear, for small changes in price it can be well approximated by a linear function. Therefore the correctly linearised model – which can be approximated by interacting price

with functions of expenditure – performs well when analysing impacts of small price changes.

The fact that, in the case of small price changes, a linear approximation of the utility model in which the marginal utility of expenditure is nonconstant succeeds in recovering the distributional patterns across consumers is potentially useful because computing counterfactual equilibria and evaluating welfare effects of a price change in the model with nonlinear utility can be considerably more costly. In particular, when expenditure enters utility nonlinearly the simple formula for compensating variation from Small and Rosen (1981) is not valid and to compute compensating variation one must use either the simulation methods introduced in McFadden (1999) and Herriges and Kling (1999) or Dagsvik and Karlström (2005). Bhattacharya (2015) shows how to estimate the marginal distribution of compensating variation non-parametrically when interest centres on the impact of a change in the price of a single good. However, in differentiated product markets in which interest typically centres on estimation of the welfare impacts of simultaneous changes in multiple prices, these methods are not applicable.

This reliance of the discrete choice literature on restrictive assumptions about the marginal utility of expenditure contrasts with the continuous choice demand literature which has concerned itself with allowing for increasingly general forms of income effects (see for instance, Deaton and Muellbauer (1980), Banks et al. (1997), Lewbel and Pendakur (2009), Hausman and Newey (2014)). Researchers in the continuous choice demand literature have found that flexible models of income effects are important for understanding demand patterns. We find that the same is true in discrete choice models.

How expenditure enters utility in random utility models may also have a strong bearing on patterns of tax pass-through and on price increases predicted by merger simulations. A series of papers (including Seade (1985), Delipalla and Keen (1992) and Anderson et al. (2001)) provide theoretical pass-through results in stylised models of imperfect competition (with either homogenous or symmetrically differentiated goods). Weyl and Fabinger (2013) provide a framework which nests many of the previous theoretical results, and highlights the importance of a number of determinants of pass-through. All of these papers highlight the important role that the curvature of market demand plays in determining tax pass-through. Constraining the marginal

utility of expenditure in logit demand models restricts the curvature of individual consumer level demand curves. Market demand curves may still be somewhat more flexible if preference heterogeneity is included in the model, but they are nonetheless influenced by assumptions made about consumer level demands. By allowing for flexibility in the marginal utility of expenditure we significantly relax these demand curvature restrictions and are therefore able to provide new robust empirical evidence of tax pass-through.

Our work is related to a large existing literature that estimates pass-through of cost shocks and taxes to prices. A set of papers use observed tax changes to estimate pass-through. These include Besley and Rosen (1999), who exploit variation in State and local sales taxes in the US and look at the impact on prices of a number of products, Delipalla and O'Donnell (2001), who analyze the incidence of cigarette taxes in several European countries and Kenkel (2005), who uses data on how the price of alcoholic beverages changed in Alaska. Results from the literature vary, but typically these papers find complete or overshifting of specific taxes, which broadly accord with our pass-through results.

A number of papers use structural models to study equilibrium pass-through. Many of these papers find that pass-through of cost shocks is incomplete (see, for instance, Goldberg and Hellerstein (2013) and Nakamura and Zerom (2010)). An important reason for incomplete pass-through of cost shocks is that often not all cost components are affected by the shock. For instance, exchange rate movements do not directly impact the cost of non-traded inputs (Goldberg and Hellerstein (2008)). In a context where firms' marginal costs are observable (in the wholesale electricity market), Fabra and Reguant (2014) find changes in marginal costs are close to fully shifted to prices. We add to this literature by studying how equilibrium tax pass-through in an imperfectly competitive market is affected by functional form assumptions that restrict the shape of market demand.

The rest of the paper is structured as follows. In Section 2 we discuss various ways of modelling the relationship between a consumer's grocery budget and their demand and the implications for measuring consumer welfare effects. In Section 3 we discuss market level demand and how assumptions made about consumer level demand influence the curvature of the market demand curve. Section 4 presents our empirical application to the butter and margarine market. We begin by discussing the market and presenting

some reduced form evidence on the relationship between total spending and demands. We then present our empirical specification and discuss identification. Next we discuss our estimates and present the counterfactual results – highlighting the importance of a flexible utility specification and comparing the effectiveness of specific and ad valorem taxes in reducing saturated fat. A final section concludes.

2 Consumer level demand

We consider a random utility model of consumer choice (see McFadden (1981)). The consumer has a budget y available to spend. The variable y may be the consumer’s income, or it may represent the total expenditure the consumer allocates to a set of goods over which preferences are weakly separable. The consumer makes a discrete choice about which alternative $j \in \{0, 1, \dots, J\}$ to purchase. In our application to the butter and margarine market j denotes butter and margarine products, y represents the consumer’s total grocery budget and the outside good is spending on all other (non butter and margarine) groceries. We denote the price of option j as p_j . Option $j = 0$ denotes the choice not to purchase any of the goods in the market, with $p_0=0$. Each option j has associated with it a vector of observable product characteristics \mathbf{x}_j and an unobservable characteristic ε_j . Utility from selecting option j is given by $U(y - p_j, \mathbf{x}_j, \varepsilon_j)$. We refer to $U(y - p_j, \mathbf{x}_j, \varepsilon_j)$ as the consumer’s conditional utility. It is the utility obtained conditional on selecting option j and spending $y - p_j$ on the outside good. In this section, we leave implicit the dependence of U on a vector of parameters θ , some of which may be random coefficients that vary across consumers. We discuss consumer heterogeneity in more detail in Sections 3 and 4.1.

The consumer indirect utility function is given by:

$$V(\mathbf{p}, y, \mathbf{x}, \boldsymbol{\varepsilon}) = \max_{j \in \{0, \dots, J\}} U(y - p_j, \mathbf{x}_j, \varepsilon_j) \quad (2.1)$$

where $\mathbf{p} = (p_1, \dots, p_J)'$, $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_J)$ and $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_J)'$. As long as the conditional utility function, $U(y - p_j, \mathbf{x}_j, \varepsilon_j)$, is continuous and non-decreasing in $y - p_j$, $V(\mathbf{p}, y, \mathbf{x}, \boldsymbol{\varepsilon})$

satisfies the properties of an indirect utility function.¹ Consumer theory does not impose further restrictions on how $y - p_j$ enters conditional utility.

To focus on the role of income effects in the most commonly used logit model, we employ the standard assumption that ε_j is additive, independent and identically distributed across alternatives and drawn from a type I extreme value distribution. As shown in McFadden and Train (2000), any discrete choice model derived from random utility maximisation has choice probabilities that can be approximated to any degree of accuracy by a mixed logit model. So this restriction does not overly constrain the scope of our analysis as long as preference heterogeneity is included in the model. An alternative is to assume ε_j is additive and is drawn from a generalised extreme value distribution, leading, for example, to a nested logit choice model.

Under the additive assumption, an individual consumer’s conditional utility is given by:

$$U(y - p_j, \mathbf{x}_j, \varepsilon_j) = \tilde{U}(y - p_j, \mathbf{x}_j) + \varepsilon_j, \quad (2.2)$$

$$\varepsilon_j \sim \text{i.i.d. type I extreme value.}$$

and the probability the consumer selects option j is given by:

$$P_j = \frac{\exp(\tilde{U}(y - p_j, \mathbf{x}_j))}{\sum_{k \in \{0, \dots, J\}} \exp(\tilde{U}(y - p_k, \mathbf{x}_k))}. \quad (2.3)$$

The bulk of the applied literature assumes that conditional utility is linear in $y - p_j$:

$$\tilde{U}(y - p_j, \mathbf{x}_j) = \alpha(y - p_j) + g(\mathbf{x}_j). \quad (2.4)$$

This means that the conditional marginal utility of y is constant. Importantly, when comparisons are made across options, y differences out of the model. Therefore, by assumption, an increase in a consumer’s grocery budget has no impact on demand for the inside products $j > 0$. To capture the fact that choice patterns commonly vary across consumers with different budgets, it is typical to include y in the model as a “preference shifter” (see, inter alia, Nevo (2001), Berry et al. (2004), Villas-Boas (2007)). For example, the parameter α may be allowed to vary linearly across consumers with

¹It is non-increasing in prices, non-decreasing in grocery budget, homogeneous of degree zero in all prices and grocery budget, quasi-convex in prices and continuous in prices and grocery budget.

y (and possibly also with other demographic variables):

$$\alpha = \alpha_0 + \alpha_1 y + \nu \tag{2.5}$$

where ν is a random coefficient. This “preference shifter” model rules out income effects at the individual level and is ad hoc. Consumer theory does not provide a theoretical explanation for why preferences should shift with y . However, this approach does allow researchers to capture, in a reduced form way, the empirical fact that spending patterns do vary cross-sectionally with total budgets.

Papers that do allow for some nonlinearity in how $y - p_j$ enters conditional utility include Berry et al. (1995), Goldberg and Verboven (2001) and Petrin (2002). These papers consider demand for large budget share product categories (automobiles and mini-vans) and specify:

$$\tilde{U}(y - p_j, \mathbf{x}_j) = \alpha \ln(y - p_j) + g(\mathbf{x}_j). \tag{2.6}$$

In this case the conditional marginal utility of y is given by $\frac{\alpha}{y - p_j}$ and is therefore inversely proportion to $y - p_j$. This specification implies that households with higher income or expenditure are less price sensitive.

In the following sections, we explore the importance of allowing for a richer relationship between y and demands. We first discuss implications for consumer welfare and for the curvature of consumer demand. We then develop an application to the market for butter and margarine and show that this is empirically important for capturing the impact of y on individual demand elasticities and welfare.

2.1 Welfare

One important use of random utility models is to compute the welfare impacts of a change in prices, product characteristics or choice sets. In industrial organisation, the focus often is on the impact on welfare of price changes (for example, due to a merger as in Nevo (2000), or due to the introduction of a tax as in Kim and Cotterill (2008)). In environmental economics, the focus is on the impact of a change in environmental amenities. In transport economics, the focus is on public investments in transport infrastructure or on taxes or subsidies that affect various modes of transport.

In the vast majority of applications of discrete choice demand models that explicitly compute consumer welfare changes researchers use the linear utility specification (as specified in equation (2.4)) including a measure of total budget or income in the model as a preference shifter (as in equation (2.5)).² In this case, measuring consumer welfare changes is relatively straightforward. In particular, the change in consumer welfare associated with a policy change is invariant to whether it is evaluated before or after the logit shocks, $\boldsymbol{\varepsilon}$, are realised, and can be computed (conditional on realisations of any random coefficients) using the formula derived by Small and Rosen (1981).

When utility is specified as a nonlinear function of $y - p_j$, consumer welfare depends on whether it is evaluated prior to or after the logit shocks are realised (McFadden (1999)). If the logit shocks represent genuine uncertainty from the consumer perspective, it may be appropriate to use an ex ante welfare criterion based on the individual consumer's *expected utility* prior to observing $\boldsymbol{\varepsilon}$. In this case, aggregate welfare is the sum of the individual expected utilities. Conversely, if there is no uncertainty for the consumer over $\boldsymbol{\varepsilon}$ but rather the logit shocks represent cross-sectional unobserved heterogeneity, then consumer welfare changes should be based on an ex post criterion based on the individual consumer's *realised utility*. In this case, aggregate welfare is the sum or average of the individual's realised utilities. We present results adopting the latter perspective, based on realised utilities. Like Herriges and Kling (1999) we find in our application that both views yield similar estimates.

Consider baseline prices \mathbf{p} and counterfactual prices \mathbf{p}' (for instance, associated with the introduction of a tax). We measure the change in consumer welfare using compensating variation – the monetary amount required to compensate the consumer post policy change that would make them indifferent to the change.³ Individual level compensating variation, cv , associated with the price change satisfies

$$V(\mathbf{p}, y, \mathbf{x}, \boldsymbol{\varepsilon}) = V(\mathbf{p}', y - cv, \mathbf{x}, \boldsymbol{\varepsilon}). \quad (2.7)$$

²Petrin (2002) is an exception. He use the log specification given by equation (2.6), and he estimates the consumer welfare effects of the introduction of minivans to the automobile market.

³The analysis is similar for equivalent variation. When conditional utility is nonlinear in $y - p_j$, the numerical values of compensating and equivalent variation will differ.

Individual cv depends on ε and therefore is a random variable from the point of view of the econometrician. From the econometrician’s perspective, aggregate welfare is the average value of cv , $CV = \mathbb{E}(cv)$.

McFadden (1999) and Herriges and Kling (1999) develop Monte Carlo Markov chain simulation methods that allow for computation of CV in the case of a nested logit model with income effects. More recently Dagsvik and Karlström (2005) have exploited duality results applied to random utility models to characterise the distribution of cv for general random utility models. Using their methods, computation of compensating variation reduces to repeated computation of a one dimensional integral. We use their results to compute CV .

3 Market level demand and pass-through

A number of papers have highlighted that the curvature of market demand is a crucial determinant of pass-through of cost shocks and taxes to consumer prices (see, *inter alia*, Seade (1985), Anderson et al. (2001) and Weyl and Fabinger (2013)). Weyl and Fabinger (2013) emphasize that, in the context of a monopolist or symmetrically differentiated single product firm oligopoly, the curvature of the log of demand is key. For instance, in the case of a single product monopolist with constant marginal costs, pass-through of a cost shock will be incomplete if and only if the monopolist faces a demand curve that is log-concave.⁴ In this case, restricting market demand to be log-concave rules out pass-through exceeding 100% by assumption. More generally, assuming a particular degree of concavity or convexity of log demand will not necessarily imply under or over-shifting exactly, but will nonetheless place strong restrictions on the possible range of pass-through. In particular, in the logit demand model, heterogeneity in consumer types and the functional form of $\tilde{U}(y - p_j, \mathbf{x}_j)$ both have a strong

⁴Let the demand curve be $q(p)$ and constant marginal cost be c . Optimisation implies $q + p \frac{dq}{dp} = c \frac{dq}{dp}$. Differentiating with respect to cost and substituting yields pass-through as

$$\frac{dp}{dc} = \frac{1}{2 - q \frac{d^2 q}{dp^2} / (\frac{dq}{dp})^2} = \frac{1}{1 - \left(\frac{d^2 \ln q}{dp^2} \right) \left(q / \frac{dq}{dp} \right)^2}.$$

This expression shows that pass-through will be incomplete ($\frac{dp}{dc} < 1$) if and only if demand is log-concave ($\frac{d^2 \ln q}{dp^2} < 0$).

bearing on the permissible curvature of the log of market demand, and therefore on pass-through.

Consider the demand curve for a product in the market. Let each consumer be indexed by (y, θ) where as discussed above y measures the consumer's income or total expenditure on a set of goods for which preferences are weakly separable (groceries in our application) and θ measures all other observable and unobservable consumer attributes that enter into utility. Normalising the size of the market to one, the market demand curve for option j is then given by:

$$q_j(\mathbf{p}) = \int P_j(y, \theta)g(y, \theta)dyd\theta, \quad (3.1)$$

where $P_j(y, \theta)$ is the individual purchase probability (in the logit case this is given by equation (2.3)) and $g(y, \theta)$ is the joint density over the elements of (y, θ) . The second derivative of the log of market demand with respect to price is given by:

$$\begin{aligned} \frac{\partial^2 \ln q_j}{\partial p_j^2} = & \int \frac{P_j(y, \theta)}{q_j} \frac{\partial^2 \ln P_j(y, \theta)}{\partial p_j^2} g(y, \theta) dyd\theta \\ & + \left[\int \frac{P_j(y, \theta)}{q_j} \left(\frac{\partial \ln P_j(y, \theta)}{\partial p_j} \right)^2 g(y, \theta) dyd\theta - \left(\int \frac{P_j(y, \theta)}{q_j} \frac{\partial \ln P_j(y, \theta)}{\partial p_j} g(y, \theta) dyd\theta \right)^2 \right]. \end{aligned} \quad (3.2)$$

The curvature of the log of market demand depends on two terms. The first term is the weighted average of the second derivatives of log individual demand, where weights are consumers' contribution to the market demand curve. The second term is the weighted variance of the slope of log individual level demand. The first term is negative if individual level demand is log-concave. The second term is non-negative and is positive when there is heterogeneity in individual demands. Log demand will be concave if individual demand is log-concave and if the cross-sectional variance of the slope of log demand is not too big. It will be convex if individual log demand is convex or if the variance term is large enough in magnitude.

In the case of a linear utility logit model with no heterogeneity, $\frac{\partial^2 \ln q_j}{\partial p_j^2}$ collapses to the second derivative of the log of individual level demand:

$$\frac{\partial^2 \ln q_j}{\partial p_j^2} = \frac{\partial^2 \ln P_j}{\partial p_j^2} = -\alpha^2 P_j(1 - P_j) < 0.$$

The curvature of the log of market demand is then fully determined by the marginal utility parameter, α , and the market share. Both individual and market demand are restricted to be log-concave. Adding heterogeneity in consumer preferences maintains the restriction on individual demand but allows for the possibility that the market demand curve might be log-convex or even be log-concave in some regions and log-convex in others.

Allowing $y - p_j$ to enter utility in a flexible nonlinear way relaxes restrictions on the curvature of both individual level and market demand. In particular, with nonlinear utility, individual level demand need not be constrained to be log-concave. The second derivative of the log of consumer demand for option j with respect to its own price is given by:

$$\frac{\partial^2 \ln P_j}{\partial p_j^2} = (1 - P_j) \left[\frac{\partial^2 \tilde{U}(y - p_j, \mathbf{x}_j)}{\partial (y - p_j)^2} - \left(\frac{\partial \tilde{U}(y - p_j, \mathbf{x}_j)}{\partial (y - p_j)} \right)^2 P_j \right]. \quad (3.3)$$

The degree of log-concavity (or convexity) is determined by the shape of the function \tilde{U} , and therefore the flexibility of the curvature of individual demand depends on the flexibility of the function \tilde{U} . If $y - p_j$ enters utility in logs as in equation (2.6), the curvature of consumer level demand is very restricted, and is log concave. However, more flexible forms of the function \tilde{U} allow for more flexibility in consumer level demand curvature including the possibility that consumer demand is log-convex in some regions (individual demand will be log-convex if \tilde{U} is sufficiently convex). Therefore, specifying utility to be a flexible nonlinear function of $y - p_j$ allows for flexibility in the curvature of market demand both through influencing the variance of the slope of individual demands and through relaxing curvature restrictions on individual demands.

4 Application

To illustrate the potential empirical importance of modelling conditional utility as a flexible nonlinear function, we provide an example using the UK market for butter and margarine. Since spending on butter and margarine accounts for just over 1% of households' grocery expenditure, this is a market in which a restriction on the marginal utility of grocery budgets might be expected to play a limited role. We

assume preferences for groceries are weakly separable, measure the grocery budget of each household using their mean weekly grocery expenditure and estimate demand under a number of different assumptions about how mean weekly grocery expenditure enters conditional utility. We compute individual and market level demand elasticities and simulate the introduction of a tax, comparing tax pass-through and consumer welfare predictions of the various specifications.

Despite being a small budget share category, butter and margarine contribute over 13% of UK households' total annual saturated fat purchases. Butter and margarine were among the set of products subject to the aborted Danish "fat tax" introduced in 2011. We provide evidence on how the structure of a tax affects its effectiveness in lowering saturated fat purchases by comparing the implications of introducing a specific tax and an ad valorem tax.

We conduct our analysis using purchase data on 10,012 households from Kantar WorldPanel for calendar year 2010. The data include details of all the purchases of groceries brought into the home that these households made throughout the year, including details of transaction prices and the nutritional content of products.

4.1 The market for butter and margarine

The UK butter and margarine market has 47 main products.⁵ There are 32 different brands. Products that have the same brand differ in terms of pack size.⁶ As we would like to allow for the possibility that households respond to a tax by down- or up-sizing product, we model demand at the brand-pack size level (in contrast to most studies, which model brand demand). The 47 products are manufactured by 8 firms. Table 4.1 lists the firms (manufacturers) that operate in the market, the brands that these firms sell, the pack sizes that each brand is available in and the products the firms sell. Unilever is the largest firm, marketing 17 products that together have a market share of 52%. The second largest is Dairy Crest with a market share of 26%, followed by Arla with 17%, and Tesco with 3%. The concentrated and asymmetric market structure will be an important determinant of pass-through of taxes to consumer prices.

⁵There are dozens of additional products with very small markets shares, ($< 0.5\%$), which we omit from our analysis.

⁶In a few instances a brand-pack size contains two products - a salted and unsalted version.

Table 4.1 also summarises the market share, mean price, butter or margarine classification, saturated fat content and mean advertising expenditure for each product. In the UK grocery market prices and advertising levels are typically set nationally.⁷ We therefore define markets temporally (in particular, monthly). In each market we compute the price of each product by calculating the mean of transaction level prices. Note that brands in the market exhibit nonlinear pricing across pack sizes. For instance, the mean price of the 250g pack of Lurpak spread is £1.42, while the mean price of the 500g pack is £2.15. The price of the larger pack is less than double. Differential time variation in within brand, cross pack size, price schedules provides a source of price variation that we exploit in estimation (and that is absent in most studies, which only model brand level demand).

There is a large degree of variation in the saturated fat content of products – ranging from 5.1g per 100g for Flora extra light (a margarine) to 54.7g per 100g for Country Life unsalted (a butter). How pass-through of a tax differs across the product saturated fat distribution will be key to determining the effectiveness of any tax in reducing saturated fat purchases.

In Table 4.2 we present some descriptive statistics on the households in our sample. The top panel contains the 10th, 25th, 50th, 75th and 90th percentiles of the distribution of mean weekly grocery expenditure over the year across households⁸. The median household spends around £41.88 per week on average on groceries and the 10th and 90th percentiles are £22.61 and £66.20 respectively. The second panel shows the percentage of households with 1, 2 up to 5 or more members. The most common household composition is 2 members. The final panel shows the labour market status of households. We distinguish between households in which all adults are in full time employment, a fraction of adults are in full time employment, no adults in full time employment and retired households. There is roughly an even split of household between fully full time, non full time and retired households, with the remaining category of a fraction of adults in full time employment accounting for just 7% of households.

⁷For instance, most supermarkets implement a national pricing policy following the Competition Commission’s investigation into supermarket behaviour (Competition Commission (2000)).

⁸By grocery expenditure we mean the household’s total expenditure on fast-moving consumer goods, these are products bought in supermarkets and taken home (including food, cleaning products and toiletries).

Table 4.1: *Products and characteristics*

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Firm	Brand	Pack size (Kg)	Product	Quantity share (%)	Price (£)	Butter or margarine	Saturated fat per 100g (g)	Advertising (£m)
<i>Adams</i>	Kerrygold	0.25	Kerrygold 250g	1.00	1.09	Butter	48.93	0.22
<i>Arla</i>				17.76				
	Anchor	0.25	Ar: Anchor 250g	1.48	1.12	Butter	54.00	0.81
	Anchor spr.	0.25	Ar: Anchor spr. 250g	0.38	1.40	Butter	31.20	0.10
	Anchor spr.	0.50	Ar: Anchor spr. 500g	3.03	1.99	Butter	31.20	0.10
	Anchor spr. lighter	0.50	Ar: Anchor light spr. 500g	1.20	1.99	Butter	23.70	0.03
	Lurpak	0.25	Ar: Lurpak ss 500g	0.82	1.17	Butter	52.00	0.73
	Lurpak	0.25	Ar: Lurpak us 250g	0.36	1.17	Butter	53.00	0.73
	Lurpak spr.	0.25	Ar: Lurpak spread ss 250g	0.59	1.42	Butter	36.70	0.04
	Lurpak spr.	0.50	Ar: Lurpak spread ss 500g	5.10	2.15	Butter	36.70	0.04
	Lurpak spr. lighter	0.50	Ar: Lurpak light ss 500g	3.48	2.17	Butter	25.80	0.00
	Lurpak spr. lighter	0.25	Ar: Lurpak light ss 250g	0.55	1.42	Butter	25.80	0.00
<i>Asda</i>	Asda	0.25	Asda 250g	0.73	0.92	Butter	54.00	0.00
<i>Dairy Crest</i>				25.39				
	Clover diet low fat spr.	0.50	DC: Clover diet 500g	1.46	1.26	Margarine	23.50	0.00
	Clover spr.	0.50	DC: Clover spr. 500g	5.21	1.20	Margarine	26.90	0.78
	Clover spr.	1.00	DC: Clover spr. 1kg	4.34	2.37	Margarine	26.90	0.78
	Country Life	0.25	DC: Country life 250g	1.16	1.08	Butter	54.00	0.47
	Country Life	0.25	DC: Country life us 250g	0.39	1.07	Butter	54.70	0.47
	Country Life light spr.	0.50	DC: Country life spr. 500g	1.08	1.87	Butter	23.00	0.00
	Country Life spr.	0.50	DC: Country life spr. 500g	1.24	2.13	Butter	31.40	0.00
	Utterly Butterly	0.50	DC: Utterly Butterly 500g	5.99	0.80	Margarine	14.70	0.15
	Utterly Butterly	1.00	DC: Utterly Butterly 1kg	1.36	1.95	Margarine	14.70	0.15
	Vitalite	0.50	DC: Vitalite 500g	2.17	0.92	Margarine	13.40	0.00
	Willow	0.25	DC: Willow 250g	0.97	0.67	Margarine	17.06	0.00
<i>Sainsburys</i>				0.91				
	Sainsburys	0.25	Sainsburys 250g	0.66	0.93	Butter	54.00	0.00
	Sainsburys	0.25	Sainsburys us 250g	0.24	0.93	Butter	54.00	0.00
<i>Morrisons</i>	Morrisons	0.25	Morrisons 250g	0.36	0.90	Butter	52.10	0.00
<i>Tesco</i>				3.15				
	Tesco butter me up	0.50	Tesco butter me up 500g	1.06	0.97	Margarine	17.50	0.00
	Tesco blended	0.25	Tesco blended 250g	0.34	1.03	Butter	48.60	0.00
	Tesco value	0.25	Tesco value 250g	1.43	0.92	Butter	48.60	0.00
	Tesco value	0.25	Tesco value us 250g	0.32	0.93	Butter	48.60	0.00
<i>Unilever</i>				51.54				
	Bertolli light olive spr.	0.50	Un: Bertolli light 500g	1.49	1.35	Margarine	9.50	0.21
	Bertolli spr.	1.00	Un: Bertolli 1kg	1.74	2.66	Margarine	14.00	1.00
	Bertolli spr.	0.50	Un: Bertolli 500g	2.59	1.34	Margarine	14.00	1.00
	Flora buttery	0.50	Un: Flora buttery 500g	7.53	1.02	Margarine	15.60	0.76
	Flora extra light spr.	0.50	Un: Flora extra light 500g	0.98	1.34	Margarine	5.10	0.60
	Flora light spr.	0.50	Un: Flora light 500g	3.55	1.22	Margarine	9.30	0.60
	Flora light spr.	1.00	Un: Flora light 1kg	4.64	2.29	Margarine	9.30	0.60
	Flora ProActiv spr.	0.25	Un: Flora proactive 250g	0.46	1.87	Margarine	8.00	1.37
	Flora ProActiv spr.	0.50	Un: Flora proactive 500g	0.82	3.62	Margarine	8.00	1.37
	Flora	0.50	Un: Flora 500g	2.37	1.22	Margarine	12.00	0.11
	Flora	1.00	Un: Flora 1kg	2.45	2.30	Margarine	12.00	0.11
	ICBINB	1.00	Un: ICBINB 1kg	2.02	2.02	Margarine	19.90	0.35
	ICBINB	0.50	Un: ICBINB 500g	7.72	0.85	Margarine	19.90	0.35
	ICBINB light	0.50	Un: ICBINB light 500g	3.70	0.87	Margarine	11.00	0.11
	Stork baking block	0.25	Un: Stork 250g	0.66	0.50	Margarine	25.70	0.19
	Stork pack	0.50	Un: Stork 500g	3.00	0.72	Margarine	14.80	0.11
	Stork pack	1.00	Un: Stork 1kg	5.82	1.34	Margarine	14.80	0.11

Notes: Price and advertising are unweighted means across all 12 markets. Quantity shares are computed using 50,060 observations in our data, weighted to reflect the British population.

Table 4.2: *Households*

Percentile	Mean weekly grocery expenditure
10th	£22.61
25th	£30.42
50th	£41.88
75th	£54.63
90th	£78.90
No. of household members	% of households
1	21.3
2	42.2
3	16.7
4	14.5
5+	6.3
Labour market status	% of households
Full time working household	30.5
One member full time work	7.3
Non full time working household	30.3
Retired household	31.8

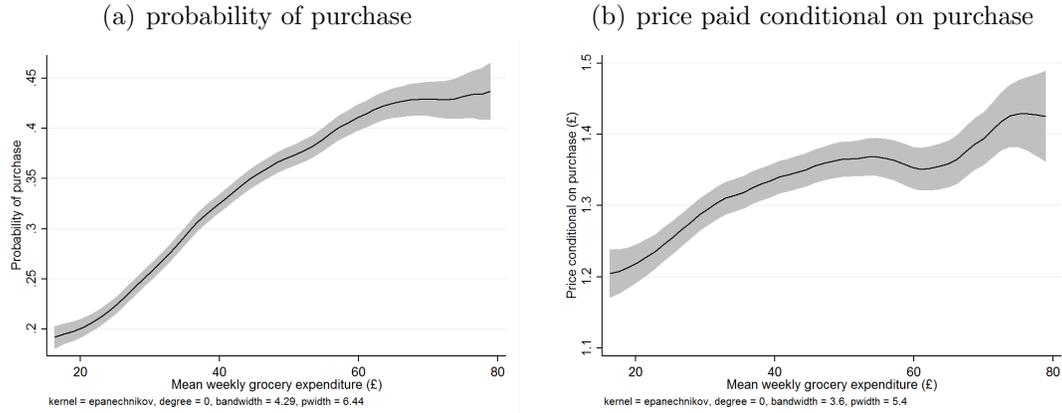
Notes: Based on our sample of 10,012 households.

4.2 Reduced form evidence

Before discussing estimates of our demand model, we provide some reduced form evidence that, despite the fact that butter and margarine are small budget share items, important features of households' butter and margarine purchasing behaviour are related to their mean weekly grocery expenditure. For each household we compute mean weekly grocery expenditure over the year, the average probability of purchasing a butter or margarine product in any given week and the mean price paid for butter or margarine conditional on purchasing. Figure 4.1 describes the relationship between these variables in the data. Panel (a) shows results of a nonparametric regression (kernel-weighted local polynomial) of the probability of purchase of a butter or margarine product on mean weekly grocery expenditure. Panel (b) shows results of a nonparametric regression of price paid conditional on purchase on mean weekly gro-

cery expenditure. The figure shows that higher mean weekly grocery expenditure is strongly correlated with both the probability of purchase, and conditional on purchase, the price of the product chosen.

Figure 4.1: *Nonparametric regression of mean weekly grocery expenditure with:*



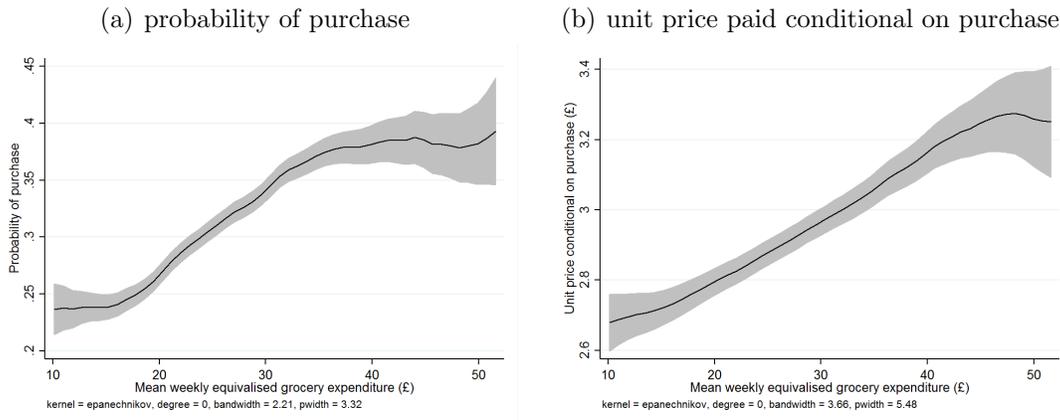
Notes: The figures display results from kernel-weighted local polynomial regressions across 10,012 households. We use weights to ensure the sample is representative of the British population. The left panel shows results from a regression of households' mean probability of purchasing butter or margarine on mean weekly grocery expenditure. The right panel shows results from a regression of households' mean price paid for butter or margarine conditional on purchase on mean weekly grocery expenditure. The shaded areas depict pointwise 95% confidence intervals.

The relationship in Figure 4.1 may reflect the causal impact of a larger grocery budget on butter and margarine demand. That is high spending households may choose higher quality (and higher priced) products because of lower price sensitivity. However, it may also reflect the effect of omitted demographics that are correlated with both grocery spending and butter and margarine purchase behaviour. For instance, it may be that larger households purchase butter or margarine more often and/or in larger pack sizes and spend more on groceries overall, but increasing the spending of a household of a given size may not influence their butter and margarine demand.

To control for the potentially confounding influence of household composition, we modify the descriptive regressions shown in Figure 4.1 in two ways. Firstly we equiv-

alize mean weekly grocery expenditure using the OECD modified scale.⁹ Secondly, we replace price paid with unit price (i.e. price per kg). These modifications should, at least in part, control for the effect of household composition and pack size purchase. Figure 4.2 shows that, even with these changes, we continue to find a positive relationship between grocery spending with both probability of purchase and (unit) price paid.

Figure 4.2: *Nonparametric regression of mean weekly equalised grocery expenditure with:*



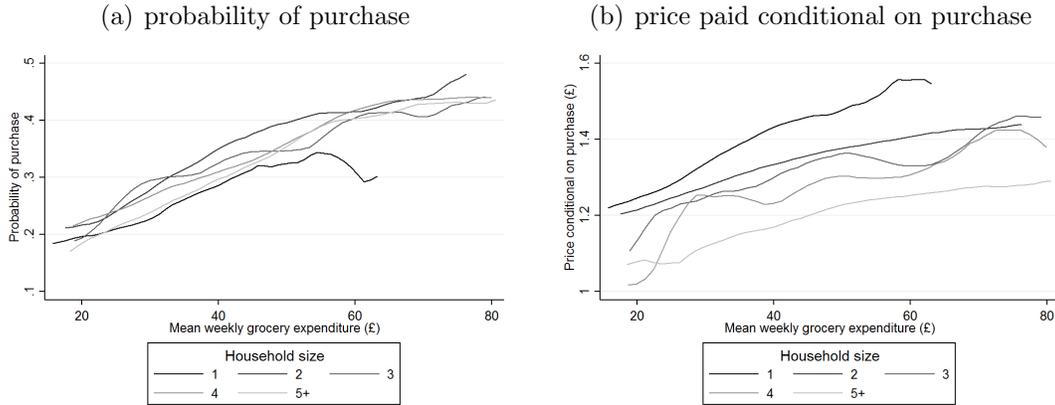
Notes: The figures display results from kernel-weighted local polynomial regressions across 10,012 households. We use weights to ensure the sample is representative of the British population. The left panel shows results from a regression of households' mean probability of purchasing butter or margarine on mean weekly equalised grocery expenditure. The right panel shows results from a regression of households' mean price per kg paid for butter or margarine conditional on purchase on mean weekly equalised grocery expenditure. The shaded areas depict pointwise 95% confidence intervals.

An alternative way to control for the influence of household composition is, instead of equalizing mean weekly grocery expenditure, simply to run the nonparametric regressions separately by household size. Panels (a) and (b) of Figure 4.3 show that for households of the same size, the relationship between a higher grocery budget and both higher purchase probability and higher price paid continues to hold. Panels (c) and

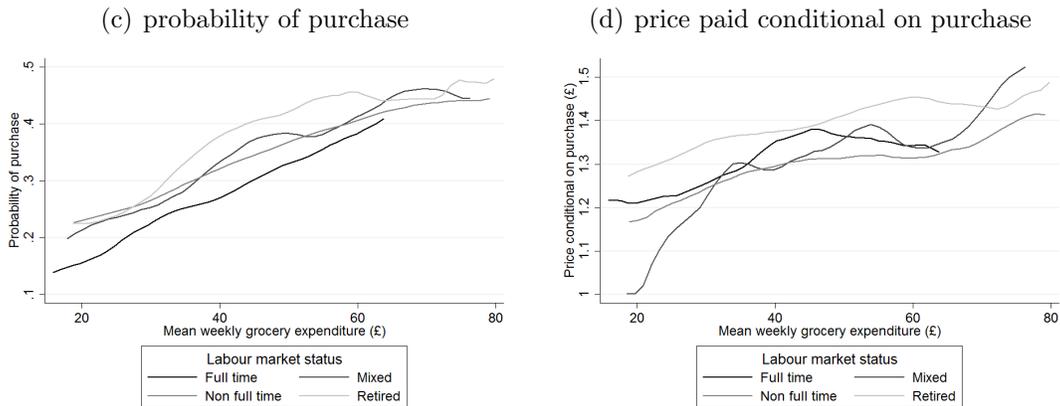
⁹This equivalence scale was introduced in Hagenaars et al. (1994) and is the most widely used scale. It assumes for every additional adult and for every person younger than 14 a household needs 0.5 and 0.3 times the resources of the first adult, respectively.

(d) show that if we condition on the labour market status of households, the pattern persists.

Figure 4.3: *Nonparametric regression of mean weekly grocery expenditure with:
By household size*



By labour market status



Notes: The figures display results from kernel-weighted local polynomial regressions across 10,012 households, ran separately by household size and labour market status. We use weights to ensure the sample is representative of the British population. The left panels shows results from regressions of households' mean probability of purchasing butter or margarine on mean weekly grocery expenditure. The right panels shows results from regressions of households' mean price paid for butter or margarine conditional on purchase on mean weekly grocery expenditure.

In summary, the descriptive evidence shows a robust relationship between mean weekly grocery expenditure and butter and margarine demand. Comparing households

that are similar in terms of their composition and labour market status, households with higher grocery budgets tend to participate in the market more frequently and purchase more expensive varieties. In what follows we allow for this possibility in a structural model of demand and supply in the butter and margarine market and highlight the importance of doing so relative to standard specifications which highly constrain the relationship.

An alternative would be to model directly the relationship between demand for butter and margarine and household income. We do not do this because we do not have good measures of household income in our data. Using different data, the Living Costs and Food Survey¹⁰, we show that the correlation between the log of equivalized household income and grocery expenditure is 0.23 and is highly statistically significant. This relationship is robust to the inclusion of expenditure on food outside the home, suggesting that income and grocery expenditure have a strong relationship, and that low levels of grocery expenditure are driven largely by poverty, not by greater reliance on or substitution to food outside the home.

4.3 Empirical model of demand

Let i index consumers and t denote time. We assume preferences for groceries are weakly separable from other goods and measure y_i as consumer i 's mean weekly grocery expenditure over the year. Let $j \in \{1, \dots, J\}$ index butter and margarine products; $j = 0$ denotes the option of not purchasing any butter or margarine product in the week. A product is a specific brand-pack size combination; we index brands by $b = 1, \dots, B$. As product characteristics, we include a vector of non-price marketing variables, \mathbf{m}_{jt} , a dummy variable for whether the product is butter, w_j , a set of pack size dummies, z_j and a brand effect, ξ_b .

We assume the utility from selecting butter or margarine product $j > 0$ takes the form:

$$U_{ijt} = f(y_i - p_{jt}; \alpha_i) + \beta' \mathbf{m}_{jt} + \gamma_i w_j + \delta'_i z_j + \xi_b + \varepsilon_{ijt} \quad (4.1)$$

and utility from not purchasing any butter or margarine is given by:

$$U_{i0t} = f(y_i; \alpha_i) + \xi_{i0} + \varepsilon_{i0t}. \quad (4.2)$$

¹⁰The Living Costs and Food Survey (LCFS) is a two week diary survey used to measure the spending patterns of UK households. We use data from 2010.

We assume that $\varepsilon_{ijt} \sim$ i.i.d. type I extreme value.

We consider three alternative forms for $f(y_i - p_{jt}; \alpha_i)$:

$$\begin{aligned} \text{Polynomial utility: } f(y_i - p_{jt}; \alpha_i) &= \sum_{n=1}^N \alpha_i^{(n)} (y_i - p_{jt})^n \\ \text{Log utility: } f(y_i - p_{jt}; \alpha_i) &= \sum_{n=1}^N \alpha_i^{(n)} \ln(y_i - p_{jt})^n \\ \text{Preference shifter: } f(y_i - p_{jt}; \alpha_i) &= \left(\sum_{n=1}^N \alpha_i^{(n)} y_i^{n-1} \right) p_{jt} \end{aligned}$$

where the coefficients are given by:

$$\begin{aligned} \alpha_i^{(1)} &= \bar{\alpha}^{(1)} + Dh_i' \alpha^{(1)} + \nu_i^\alpha \\ \alpha_i^{(n)} &= \bar{\alpha}^{(n)} + Dh_i' \alpha^{(n)} \quad \text{for } n > 1. \end{aligned}$$

Dh_i denotes a set of dummy variables capturing household size. We interact all of the α_i parameters with household size, meaning we allow the function $f(y_i - p_{jt}; \alpha_i)$ to vary across households of different sizes. We also include a random coefficient on the first order term in $f(y_i - p_{jt}; \alpha_i)$, capturing unobserved differences across households.

Both the polynomial utility (as long as $n > 1$) and the log utility models allow for a nonlinear relationship between grocery expenditure and demand. The preference shifter model is not designed to capture a causal effect of grocery expenditure on demand; rather grocery expenditure is included in an ad hoc way to capture cross-sectional correlation in grocery expenditure and demands. The two standard specifications in the existing literature correspond to the preference shifter model with $n = 2$ or the log utility model with $n = 1$.

We allow for preference heterogeneity for the product characteristics. In particular, we model the preference for the attribute ‘‘butter’’ (γ_i) as a random coefficient. We assume the $f(y_i - p_{jt}; \alpha_i)$ and butter random coefficients are normally distributed (and potentially correlated). To capture the fact that large households may prefer large pack sizes we allow the coefficients on the pack size dummies (δ_i) to shift linearly with household size. To capture the fact that participation in the market may vary

across households of different sizes and with different labour market status we allow the baseline utility from the no purchase option (ξ_{i0}) to vary across these dimensions.

The non-price marketing variables we control for include brand advertising and product promotions. In particular, we observe monthly expenditure on advertising of each brand. To control for the potentially dynamic effects of advertising (see, for instance, Dubé et al. (2005)), we compute a stock of brand advertising and control for this and its square (thereby allowing for the possibility of diminishing effects of advertising on utility). We also observe whether a product is on promotion. Being on promotion is clearly correlated with lower prices but also captures marketing efforts over and above price cuts and national advertising expenditures. The inclusion of controls for these marketing activities captures important time varying efforts that firms make to shift demand for their products.

4.4 Identification

The principal identification issue is whether we are able to correctly identify the relationship governing how grocery expenditure and price affect demand. There are four main threats to identification.

Firstly, it is possible that there exist shocks to butter and margarine demand that are correlated with price. The most obvious factors that might lead to unobserved changes in demand that are correlated with prices are unobserved advertising or promotional efforts or time-varying brand quality effects. We directly control for both advertising and promotional activity. This limits the possibility that unobserved advertising or promotions are shifting both demand and price around over time.

We directly control for both brand and pack size fixed effects. We can still identify the impact of price (and grocery expenditure) on demand because we can exploit two forms of price variation (and variation in grocery expenditure) in the data. First, prices vary over time conditional on brand, pack size and time-varying promotional activity. Second, prices vary within brand across pack sizes, conditional on brand and pack size fixed effects. This is due to nonlinear pricing across pack sizes that is common in the UK (prices are linear for a fixed pack size but nonlinear across different pack sizes of the same brand). This price variation is not collinear with the pack size fixed effects. In addition, the extent of nonlinear pricing varies over time. We think that it is unlikely that there are systematic shocks to households' valuation of pack sizes that are

differential across brand once we condition on brand and pack size fixed effects as well as on advertising and promotional activity. Rather, it is more plausible that such tilting of brand price schedules is driven by cost variations that are not proportional to pack size, differential pass-through of cost shocks and differences in how brand advertising affects demands for different pack sizes. This identification argument is similar to that in Bajari and Benkard (2005). In an application to the computer market, they assume that, conditional on observables, unobserved product characteristics are the same for all products that belong to the same model. We assume that unobserved product characteristics do not vary within brand (conditional on advertising and promotions).

The second main threat to identification is the possible existence of shocks to butter and margarine demand that are correlated with grocery expenditure. To avoid endogeneity concerns about trip-level grocery expenditure we measure household expenditure as the household's mean weekly grocery expenditure over the year. If we were to measure grocery expenditure at the shopping trip level, a concern might be that trip level expenditure is correlated with idiosyncratic shocks to butter and margarine demand (a birthday party or festivity leading to a positive demand shock for butter and higher grocery expenditure on that shopping trip). A second issue might be that much of the high frequency variation in trip level expenditure is likely to reflect planning decisions related to how many shopping trips the household plans to undertake in a given period of time, and would not be informative about income effects. The use of mean weekly grocery expenditure over the year minimises these concerns and ensures that we only exploit variation in grocery budgets that reflect long run expenditure decisions.

The third main threat to identification is that variation in omitted demographic characteristics might drive both differences in grocery expenditure and differences in demand. The most plausible demographic driver of both butter demand and grocery expenditure is differences in household size (e.g. larger households may have larger grocery budgets and purchase more butter or margarine). We account for this both by allowing household size to directly affect households' valuation of the pack size dummies and the option to not purchase, and by allowing the function relating mean weekly grocery expenditure to utility to vary with household size. Thus, we exploit cross-sectional variation in mean weekly grocery expenditure and demand across households of the same size. As a result, omitted household size effects are not driving our results.

A second plausible driver of butter and margarine demand is labour force status. We assume preferences for grocery demand are weakly separable from non-grocery expenditure. In particular, we assume that decisions over non-grocery consumption and labour supply influence the choice of butter or margarine (and consumption of the “other groceries” outside good) only through their influence on grocery expenditure. This assumption might be violated if for example, retired people spend less on groceries and purchase butter or margarine more frequently. To deal with this possibility we follow Browning and Meghir (1991) and control for labour market status (allowing it to affect the valuation of the no purchase option). This controls for the possibility of some form of nonseparabilities between labour supply and grocery choice and eliminates labour force status as a source of omitted variable bias.

Finally, our main specification assumes that total grocery expenditures are independent of consumer random coefficients. This assumption enables us to estimate a flexible discrete choice model that (i) captures variation in the effect of policy reform across the total grocery expenditure distribution, (ii) that relaxes curvature restrictions on demand which impacts on tax pass-through and (iii) that includes total expenditure in a theoretically coherent way. An alternative specification that we consider is the “preference shifter” specification that includes expenditure in a flexible reduced form way as a variable that shifts the price coefficient.

The difference between these two approaches is that our preferred specification allows for flexible income effects in a theoretically coherent way, but assumes that there is no correlation between unobserved preferences for butter and margarine and total grocery expenditure. The alternative flexible preference shifter model assumes that there are no income effects, so is much more theoretically restrictive, but allows expenditure to shift the mean of the random coefficients distribution.

In our empirical application these two models yield similar predictions. We prefer the former specification because expenditure enters the model in a way that is coherent with economic theory, and we do not have to impose economically substantive restrictions on income effects.

4.5 Firm competition

Let $f = \{1, \dots, F\}$ index firms and F_f denote the set of products owned by firm f . We assume that manufacturers compete by simultaneously setting prices in a Nash-

Bertrand game. We consider a mature market with a relatively stable set of products, and we therefore abstract from entry and exit of firms and products from the market. We deploy the commonly used approach of using our demand estimates and an equilibrium pricing condition to infer firms' marginal costs (see Berry (1994) or Nevo (2001)).

In our application, we assume that the 8 manufacturing firms compete by setting prices (for example, as would be the case if retailers charge an exogenous constant markup). We do not have information about upstream bargaining and contracting between retailers and wholesalers in this market and leave exploration of the impact of this restriction to future work.

Normalising the size of the market to be one and integrating across consumers, the demand model specified in equations (4.1) and (4.2) implies a market demand function $q_j(\mathbf{p}_t)$ for each j . Then, firm f 's (variable) profits in market t are given by:

$$\Pi_{ft}(\mathbf{p}_t) = \sum_{j \in F_f} (p_{jt} - c_{jt}) q_j(\mathbf{p}_t). \quad (4.3)$$

The first order conditions for firm f are

$$q_j(\mathbf{p}_t) + \sum_{k \in F_f} (p_{kt} - c_{kt}) \frac{\partial q_k(\mathbf{p}_t)}{\partial p_{jt}} = 0 \quad \forall j \in F_f. \quad (4.4)$$

In a Nash equilibrium, the first order conditions (4.4) are satisfied for all firms. Under the assumption that observed market prices are an equilibrium outcome of the Nash-Bertrand game played by firms, given our estimates of the demand function, we can invert firms' first order conditions to infer marginal costs.

4.6 Counterfactual

Using our demand estimates and the above model of firm competition, we simulate the introduction of a specific tax (t) that is proportional to the saturated fat content of a product. Let η_j denote the saturated fat content of product j and $\boldsymbol{\eta} = (\eta_1, \dots, \eta_J)'$. A

counterfactual equilibrium price vector \mathbf{p}_t^s satisfies:

$$q_j(\mathbf{p}_t^s + t\boldsymbol{\eta}) + \sum_{k \in F_f} (p_{kt}^s - c_{kt}) \frac{\partial q_k(\mathbf{p}_t^s + t\boldsymbol{\eta})}{\partial p_{jt}} = 0 \quad \forall j \in F_f, \text{ and } \forall f \in 1, \dots, F. \quad (4.5)$$

We also consider the introduction of an ad valorem tax (τ), such that a counterfactual equilibrium price vector \mathbf{p}_t^{av} satisfies:¹¹

$$q_j((1 + \tau\boldsymbol{\eta})\mathbf{p}_t^{av}) + \sum_{k \in F_f} (p_{kt}^{av} - c_{kt}) \frac{\partial q_k((1 + \tau\boldsymbol{\eta})\mathbf{p}_t^{av})}{\partial p_{jt}} = 0 \quad \forall j \in F_f, \text{ and } \forall f \in 1, \dots, F. \quad (4.6)$$

4.7 Estimates

We estimate specifications corresponding to the polynomial utility, log utility and preference shifter models outlined in Section 4.3 using maximum likelihood. To compute the likelihood function, we use Gauss-Hermite quadrature rules to integrate out the random coefficients. For the polynomial utility model we estimate separate specifications allowing utility to be a first order, second order, up to eleventh order polynomial. The specification that minimises the Bayesian Information Criterion is the third order, cubic utility model. In Table 4.3 we report the estimated coefficients of the cubic utility model.¹²

The top panel presents estimates of the parameters of the distribution of random coefficients. We allow the cubic function relating spending on the outside good, $y - p_j$, to consumer utility to vary across households with 1, 2, 3, 4 and 5+ members. We model the first order term as a random coefficient, allowing for unobserved preference heterogeneity across households. We also include a random coefficient on the attribute “butter” (as “butter” is collinear with the brand effects, we constrain it to have zero mean). As noted above, we assume that the random coefficients are jointly normally distributed and allow for correlation between the coefficients. Direct interpretation

¹¹For both taxes theoretical conditions for the uniqueness of an equilibrium are not known. In each case we choose 500 randomly drawn starting values for the nonlinear equation solver and in each case the solver converges to the same equilibrium.

¹²In the Appendix we report coefficients for the linear utility, preference shifter and log utility specifications. We do not report coefficient estimates for all the polynomial specifications. These are available from the authors upon request.

of the coefficients of the polynomial in $y - p_j$ is not very informative. We simply note that, for each household size category, the first, second and third order terms are statistically significant as are the random coefficient variance and covariance terms. Below we analyse the implications of these estimates for marginal utility, elasticities of demand, consumer welfare and pass-through.

The bottom section of the table shows the coefficient estimates for the non-random coefficients. The coefficient on the promotion dummy is positive and statistically significant. This suggests that holding price and advertising fixed, other non-price marketing activities significantly increase demand. Advertising has a positive, but diminishing effect on utility, although this effect is only marginally statistically significant. We interact pack size effects with household size and find larger households are more likely to purchase large pack sizes. The interactions of household size and labour market status with the no purchase option indicate that larger households and households with two full time working adults purchase butter or margarine more frequently.

Table 4.3: *Coefficient estimates: cubic utility*

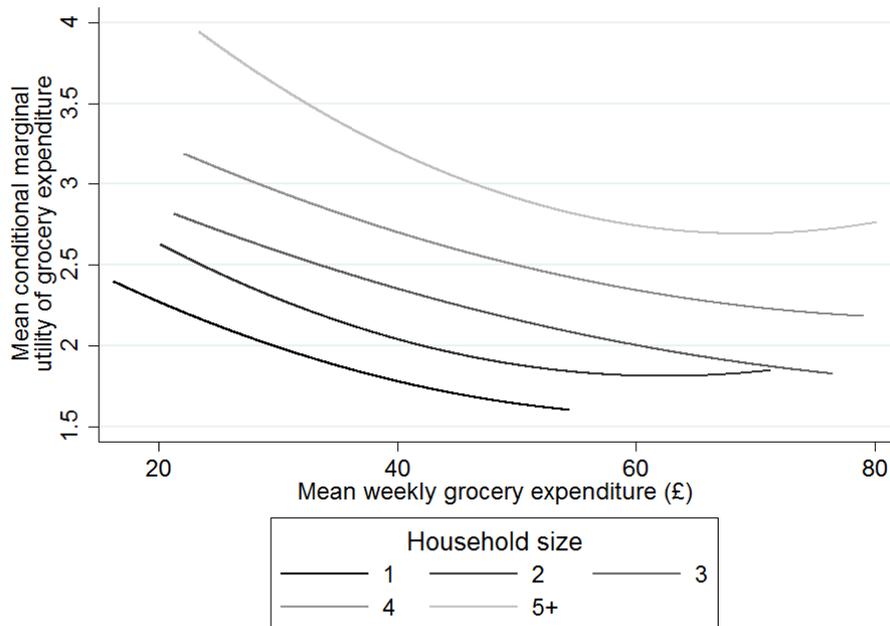
	Coefficient estimate	Standard error
Random coefficients		
<i>Mean terms</i>		
(y-p)	30.4837	1.8375
(y-p) ²	-2.2973	0.5097
(y-p) ³	0.1188	0.0475
<i>Interaction terms</i>		
(y-p) *hh2	5.3807	2.1565
(y-p) ² *hh2	-0.5455	0.5880
(y-p) ³ *hh2	0.0336	0.0520
(y-p) *hh3	4.5022	2.8780
(y-p) ² *hh3	0.4946	0.6938
(y-p) ³ *hh3	-0.0568	0.0567
(y-p) *hh4	9.5215	3.2152
(y-p) ² *hh4	0.1886	0.7381
(y-p) ³ *hh4	-0.0374	0.0582
(y-p) *hh5	24.8701	4.5240
(y-p) ² *hh5	-1.8109	0.9647
(y-p) ³ *hh5	0.0787	0.0702
<i>Variance-covariance terms</i>		
Var(y-p)	72.0934	1.5038
Var(Butter)	14.6233	0.4743
Cov(y-p, Butter)	1.6737	0.0337
Fixed coefficients		
Promotion	0.1924	0.0271
Advertising	0.0079	0.0052
Advertising ²	-0.0002	0.0001
500g	2.3917	0.0834
1kg	2.2948	0.1459
500g*Household size	0.2012	0.0167
1kg*Household size	0.5783	0.0292
No purchase*Mixed work	-0.2166	0.0482
No purchase*Non full time	-0.2697	0.0309
No purchase*Retired	-0.4457	0.0324
No purchase*hh2	-0.4841	0.0512
No purchase*hh3	-0.6233	0.0621
No purchase*hh4	-0.8324	0.0667
No purchase*hh5	-1.1058	0.0860
Brand fixed effects		
	Yes	
Likelihood	-8.8729e+04	
B.I.C.	1.7815e+05	

Notes: Sample size is 10,012 households with 5 observations per household. Random coefficients are assumed to be distributed joint normally. The butter dummy is collinear with the brand effects and therefore has a mean coefficient that is constrained to be zero.

In Figure 4.4 we plot the mean conditional marginal utility of grocery expenditure for each household size group. For each household size, over most of the domain, the marginal utility of grocery expenditure is a decreasing, convex function of $y - p_j$. The shape of the relationship is broadly similar across household sizes. The mechanism is as follows. When a household's grocery expenditure rises, their marginal utility of expenditure falls. As a result, high expenditure households are less price sensitive than low expenditure households. In the context of the butter and margarine market this implies that high expenditure households purchase butter and margarine more frequently and are more likely to purchase high quality products (that have higher prices). The main impact of household size is to shift the relationship up. Therefore, conditional on household size, lower levels of grocery expenditure are associated with a higher marginal utility, but conditional on grocery expenditure larger households tend to have higher marginal utilities.

In Section 3 we highlighted that allowing utility to depend on $y - p_j$ through a nonlinear function, $\tilde{U}(\cdot)$, allows for the possibility of household level demands that are log-convex (something that is typically ruled out in applied applications). Log-convex demand arises if $\tilde{U}(\cdot)$ is sufficiently convex, which requires the conditional marginal utility to be an increasing function of $y - p_j$. Figure 4.4 makes clear that in our application we do not find evidence of log-convex household demands. This result is not driven by functional form assumptions.

Figure 4.4: *Conditional marginal utility of grocery expenditure: cubic utility*



Notes: Lines shows mean conditional marginal utility of grocery expenditure for households with 1, 2, 3, 4 and 5+ members.

In Figure 4.5 we compare estimates of the mean conditional marginal utility of grocery expenditure across different model specifications. We focus on comparisons of the relationship for households with 2 members.¹³ Panel (a) shows the relationship for the linear, cubic, quartic and quintic utility specifications. The linear utility model constrains the marginal utility of grocery expenditure to be constant. This restriction is clearly not supported by the data. The quartic and quintic specifications yield estimates almost indistinguishable from the cubic specification.

Like the linear utility specification, the preference shifter specification imposes that the conditional marginal utility of grocery expenditure is constant for a given household. However, it does allow the parameter to shift linearly across households based on their grocery budget, y . Panel (b) of Figure 4.5 shows that the specification does, to some extent, capture the fact that households with higher grocery budgets have a

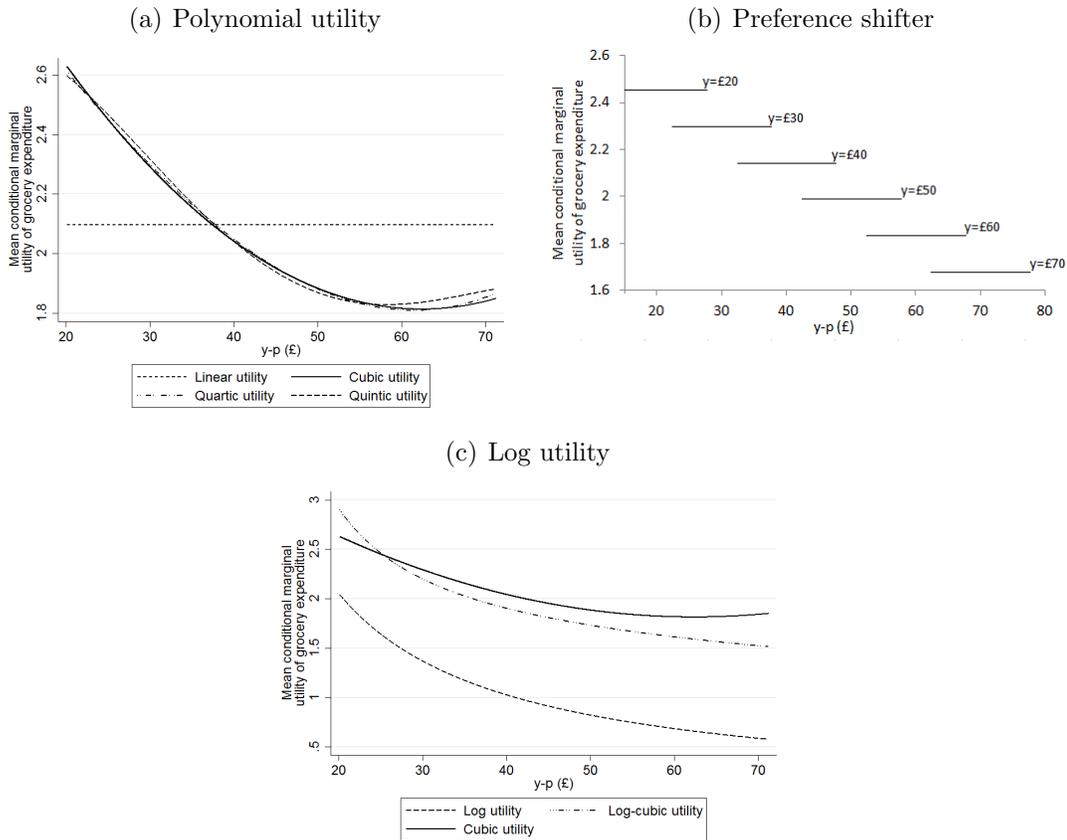
¹³We focus on households of one particular size to avoid confounding mean weekly grocery expenditure and household size effects. As households with 2 members are most numerous in our data we choose this group. The conclusions we draw about the different specifications hold equally for households of any other size.

lower marginal utility. However, the linear way in which y interacts with the coefficient on price, means the specification is unable to capture the convexity exhibited in the estimates of the cubic specification.

The log utility specification, shown on panel (c), yields an estimate of the conditional marginal utility of grocery expenditure that decreases convexly, but the function is shifted vertically downwards compared to the function implied by the cubic specification (also shown on the graph). In principle this could reflect mis-specification of the cubic utility model, or mis-specification of the log utility specification. The latter is much more likely, because specifying utility to be linear in the log of $y - p_j$ leaves only one parameter to determine the location, slope and curvature of the conditional marginal utility function. To test whether this is indeed the case we also estimate the model specifying utility as a *third order polynomial in the log of $y - p_j$* (denoted log-cubic utility in the figure). This specification, which is more general and nests the log utility specification, yields an estimate of the marginal utility that is similar to the cubic utility specification.

As our baseline model we proceed with the cubic utility specification. It is clear from panel (c) of Figure 4.5 that, in our application, the log utility specification does a poor job of replicating the shape of the conditional marginal utility of grocery expenditure found with more flexible specifications. In addition, the log utility model yields implausible estimates of marginal costs and welfare. In what follows we therefore compare our baseline model to the linear utility and preference shifter specifications.

Figure 4.5: *Conditional marginal utility of grocery expenditure for household size 2: alternative specifications*



Notes: Lines shows mean conditional marginal utility of grocery expenditure for households with two members for various utility specifications.

The market level price elasticities are crucial determinants of equilibrium prices in models of firm pricing in imperfectly competitive markets. It turns out in our empirical application that the cubic utility, linear utility and preference shifter models all yield market level price elasticity and marginal cost estimates that are very similar.¹⁴ In other words, all three specifications agree on the slope of market demand at observed prices. This need not be true in general.

While market elasticities determine the nature of the pricing equilibrium, household level elasticities are important for determining the distributional impact of a policy reform. We find in our application that, unlike the market elasticities, the household

¹⁴See the Appendix.

level elasticities are sensitive to whether we incorporate grocery expenditure in the model in a flexible and theoretically consistent way or not. To illustrate this, we compute each household's own-price elasticity of demand for butter and margarine for each choice occasion in our data (this is the market share weighted average of household's own-price elasticities across products).

In Table 4.4 we report the mean household level own price elasticity under each specification. We also report the average deviation from the mean own price elasticity for households in each quartile of the mean weekly grocery expenditure distribution. To avoid confounding expenditure effects with the influence of household size, we show the results for households with 2 members (very similar patterns holds for households of other sizes). The table also contains 95% confidence intervals.¹⁵ In Figure 4.6 we plot household level own price elasticities versus mean weekly grocery expenditure for each of the model specifications. The mean household own price elasticity is essentially the same under each model specification. However the three specifications yield different predictions for how price sensitivity varies across the grocery expenditure distribution. The cubic utility specification results indicate that households with low grocery budgets are the most price sensitive; households in the bottom quartile of the mean weekly grocery expenditure distribution, on average, have an own price elasticity 0.29 below the mean and households in the top quartile, on average, have an own price elasticity 0.21 above the mean. By assumption, the linear utility model completely fails to capture the variation in price sensitivity across the grocery expenditure distribution. The preference shifter specification does predict falling price sensitivity across the grocery expenditure distribution, but it fails to capture the concavity in the relationship, underestimating price sensitivity at the bottom and top of the distribution and overestimating it in the centre.

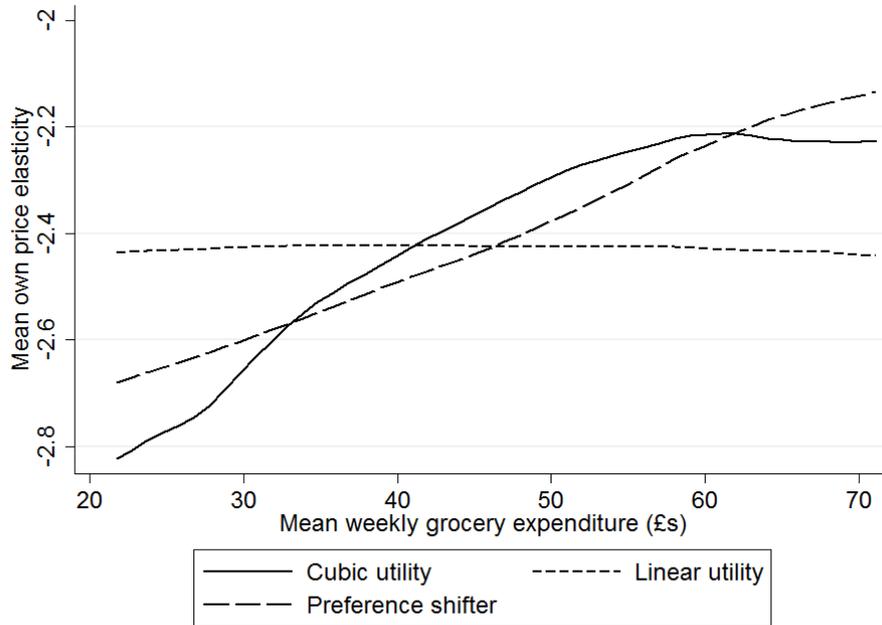
¹⁵We calculate confidence intervals in the following way. We obtain the variance-covariance matrix for the parameter vector estimates using standard asymptotic results. We then take 100 draws of the parameter vector from the joint normal asymptotic distribution of the parameters and, for each draw, compute the statistic of interest, using the resulting distribution across draws to compute Monte Carlo confidence intervals (which need not be symmetric around the statistic estimates).

Table 4.4: *Household own price elasticity for household size 2*

Specification	Mean own price elasticity	Average deviation from mean own price elasticity for quartile of grocery expenditure distribution:			
		1	2	3	4
Cubic utility	-2.44 [-2.59, -2.29]	-0.29 [-0.33, -0.26]	-0.03 [-0.05, -0.02]	0.11 [0.08, 0.13]	0.21 [0.18, 0.25]
Linear utility	-2.43 [-2.58, -2.32]	-0.01 [-0.01, 0.00]	0.01 [0.01, 0.01]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]
Preference shifter	-2.44 [-2.61, -2.29]	-0.22 [-0.25, -0.19]	-0.07 [-0.08, -0.06]	0.04 [0.03, 0.04]	0.26 [0.22, 0.29]

Notes: Numbers are for households with 2 members. For each choice occasion we compute the market-share weighted mean own price elasticity. Numbers shows average of this own price elasticity. We measure grocery expenditure as the households' mean weekly grocery expenditure over the year. 95% confidence intervals are shown in brackets.

Figure 4.6: *Variation in own price elasticities with grocery expenditure for household size 2*



Notes: Numbers are for households with 2 members. For each choice occasion we compute the market-share weighted mean own price elasticity. Figure shows local polynomial regression of how mean choice occasion elasticity varies with households' mean weekly grocery expenditure over the year.

4.8 Counterfactual results

For our counterfactual experiments, we simulate the introduction of a tax levied on the saturated fat in butter and margarine. We separately simulate the introduction of a specific and an ad valorem tax. In each case we select the rate of tax that results in a 20% fall in purchases of saturated fat in the case of no firm pricing response (i.e. in the case of 100% pass-through). For the cubic utility model, in the case of the specific tax this implies a tax rate of 13.8 pence per 100g of saturated fat and in the case of the ad valorem tax it implies a tax rate of 9.4 pence per 100g of saturated fat. The response we estimate is the short-run pricing response rather than long-run effect. In the long-run there may be increased consumer awareness of the dangers of excess saturated fat consumption, product reformulation and product entry and exit (see Anderson et al. (2001), Hamilton (2009), Draganska et al. (2009)).

We firstly focus on how assumptions about the marginal utility of grocery expenditure affect conclusions about the impact on market equilibria and the welfare effects of tax. To do this we focus on the specific tax and compare predictions across model specifications.¹⁶ We then use the cubic utility specification to compare the performance of the specific and ad valorem taxes in lowering saturated fat purchases.

4.8.1 Comparison of specifications

Figure 4.7 is a scatter plot, at the product level, that shows the relation between pass-through of the specific tax and a product’s total saturated fat content. We plot the numbers for the cubic utility specification and for three alternative specifications – the linear utility and preference shifter specifications and a multinomial logit specification with linear utility and with no consumer level heterogeneity.

For the polynomial utility specification, across all products in the market, average pass-through of the tax to consumer prices is 102%. Therefore, on average the model predicts that prices will move close to one-to-one with the specific tax. This average masks a considerable degree of heterogeneity across products. Figure 4.7 shows that products with higher saturated fat content tend to have higher tax pass-through. As the tax is levied on saturated fat content, this implies that firms’ equilibrium pricing response acts to amplify the price differential the tax creates between low and high fat products.

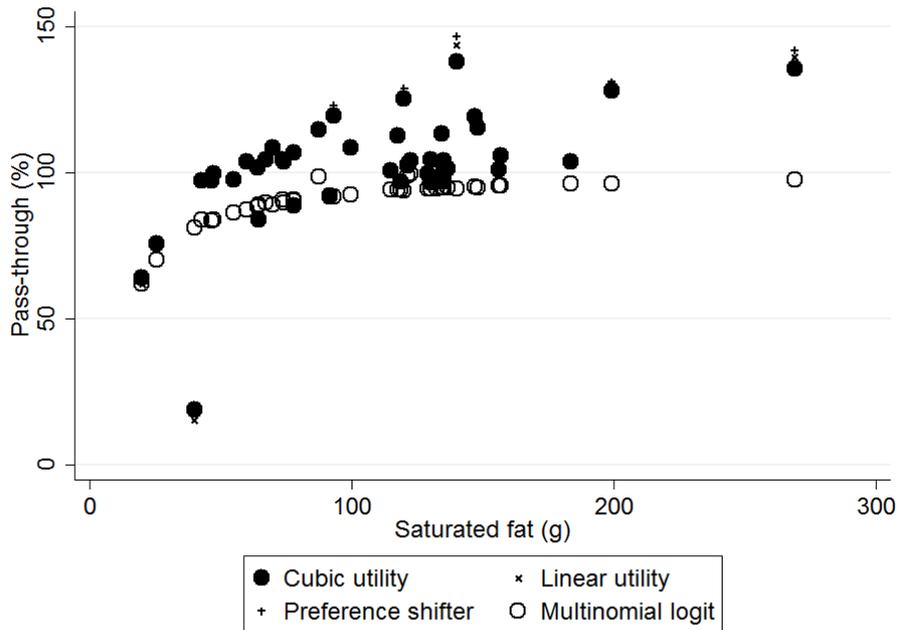
In Section 3 we highlighted that an important determinant of tax pass-through is the curvature of the log of market demand, and that an advantage of a model in which utility is flexible and nonlinear in $y - p_j$ is that it relaxes restrictions on the curvature of log market demand through allowing for more flexibility in the curvature of log household demand. This flexibility allows one to assess empirically whether the imposition of log-concave market demand is driving results on pass-through. Figure 4.7 shows that, in our application, the alternative more restrictive linear utility specification yields pass-through results that are very similar to those found by the cubic utility specification. This is also true for the preference shifter model. In this market, this suggests that the pass-through results are not driven by the curvature restrictions placed on household level demands (e.g. log-concavity) when utility is linear in $y - p_j$.

¹⁶In the Appendix we repeat the analysis for the ad valorem tax and show our conclusions about the different specifications remain unchanged.

If we had only estimated the linear utility or preference shifter models, we would not be able to provide empirical evidence on this question because we would have imposed *a priori* log-concavity.

A second determinant of the curvature of log market demand is the average variance of the slope of the log of household demand curves. In each of the cubic utility, linear utility and preference shifter specifications we allow for the possibility that the variance is non-zero through the inclusion of observed and unobserved preference heterogeneity (through household size and random coefficients). In addition, the preference shifter and cubic models also allow for positive variance through the inclusion of grocery expenditure (as a preference shifter in the first case, and as an argument of consumer level utility in the second). Allowing for this heterogeneity is important in practice. Figure 4.7 shows that a multinomial logit model that excludes any preference heterogeneity, and in which utility is specified to be linear in $y - p_j$, yields pass-through which is lower than the random coefficient models; pass-through is 92% on average. It is well known that inclusion of rich preference heterogeneity in logit demand models is important for capturing realistic substitution patterns. Our results suggest, not surprisingly, it is also important when modelling pass-through.

Figure 4.7: *Pass-through of specific tax across products*



Notes: For each product in each market we compute the pass-through of the tax. Figure is a scatter plot of products' mean pass-through across markets with their saturated fat contents.

In the first column of Table 4.5 we report average compensating variation estimated using each model specification. Again, to avoid confounding expenditure effects with the influence of household size, we show the results for households with two members (very similar patterns holds for households with other sizes). These numbers can be interpreted as the monetary payment (per year) the average 2 member household would require to be indifferent to the change in tax policy. All three models predict average compensating variation of around £2.50.

Columns two to five of Table 4.5 show the average deviation from mean compensating variation for households in each quartile of the mean weekly grocery expenditure distribution. Figure 4.8 shows graphically how compensating variation varies with mean weekly grocery expenditure. The linear utility specifications predicts no relationship. The other two specifications suggest compensating variation is increasing in grocery expenditure. The cubic utility model suggests that on average households in the bottom quartile of the mean weekly grocery expenditure distribution have compensating variation of £0.70 below average and household in the top quartile, on average,

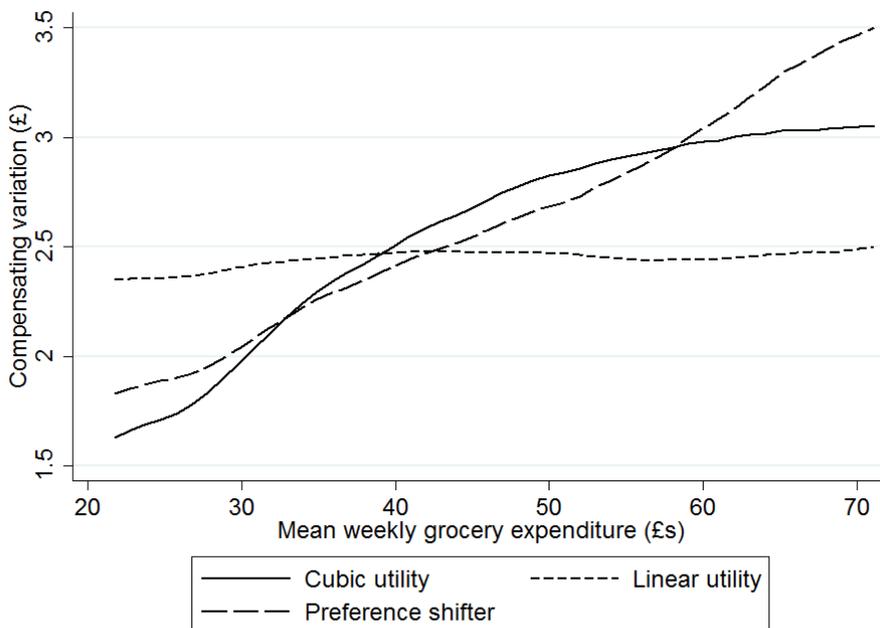
have compensating variation £0.49 above average. Households towards the bottom of the grocery expenditure distribution both purchase less butter and margarine and are more willing to switch between alternatives in response to a price change. As a result, in absolute terms they are harmed less than high expenditure households. The preference shifter model also predicts a positive relationship between a household’s mean weekly grocery expenditure and compensating variation. However, it fails to capture the concavity of the relationship and overestimates compensating variation at the bottom and top of the grocery expenditure distribution and underestimates it towards the centre.

Table 4.5: *Compensating variation from specific tax for household size 2*

Specification	Mean compensating variation	Average deviation from mean compensating variation for quartile of grocery expenditure distribution:			
		1	2	3	4
Cubic utility	2.49 [2.34, 2.69]	-0.70 [-0.80, -0.64]	-0.06 [-0.13, -0.03]	0.28 [0.23, 0.35]	0.49 [0.39, 0.61]
Linear utility	2.45 [2.32, 2.66]	-0.07 [-0.08, -0.06]	0.02 [0.01, 0.02]	0.04 [0.04, 0.05]	0.01 [0.00, 0.02]
Preference shifter	2.52 [2.35, 2.74]	-0.59 [-0.68, -0.50]	-0.18 [-0.27, -0.14]	0.10 [0.07, 0.13]	0.67 [0.54, 0.80]

Notes: Numbers are for households with 2 members. Numbers give compensating variation for the average household associated with the simulated specific tax. We measure expenditure as the households’ mean weekly grocery expenditure over the year. Numbers are for a calendar year. 95% confidence intervals are shown in brackets.

Figure 4.8: *Variation in compensating variation from specific tax with grocery expenditure for household size 2*



Notes: Numbers are for households with 2 members. Figure shows local polynomial regression of how compensating variation from specific tax varies with households' mean weekly grocery expenditure over the year.

The distributional results from the preference shifter model differ from those from the cubic utility specification because the preference shifter model does not allow enough flexibility in the way in which y enters to fully recover the variation in purchase patterns with grocery expenditure.

It is straightforward to demonstrate this empirically, and at the same time suggest a modification to the preference shifter model that can recover the full distributional consequences of the saturated fat tax. If the true model is cubic as our results suggest, a first order approximation around $p_j = 0$ is given by:

$$U_j \approx -(a^{(1)} + a^{(2)}y + a^{(3)}y^2)p_j + g(\mathbf{x}_j) + \epsilon_j \quad (4.7)$$

where we have omitted all terms that do not vary across j and where $a^{(1)} = \alpha^{(1)}$, $a^{(2)} = 2\alpha^{(2)}$ and $a^{(3)} = 3\alpha^{(3)}$. The approximation error is quadratic in p_j and depends on \tilde{U}'' . If for a given consumer the conditional marginal utility of grocery expenditure

is approximately constant in the region $[y - p_j, y]$ then the approximation will work well. When utility is smooth, this will be the case when p_j is small relative to y . In our application, estimation of the linearised utility model associated with equation (4.7) yields results, including distributional effects, which are very close to those from the cubic polynomial utility specification. While this model is not as appealing from a theoretical point of view, it may offer a practically expedient way to capture variation across the grocery expenditure or income distribution. A researcher who did not know the correct functional form for \tilde{U} could allow the price coefficient to be a nonparametric function of y . Tables 4.6 and 4.7 illustrate for both the household level elasticities and compensating variation.

Table 4.6: *Mean own price elasticity for household size 2: Cubic and linearised utility*

Specification	Mean own price elasticity	Average deviation from mean own price elasticity for quartile of grocery expenditure distribution:			
		1	2	3	4
Cubic utility	-2.44	-0.29	-0.03	0.11	0.21
	[-2.59, -2.29]	[-0.33, -0.26]	[-0.05, -0.02]	[0.08, 0.13]	[0.18, 0.25]
Linearised utility	-2.45	-0.29	-0.03	0.11	0.20
	[-2.61, -2.29]	[-0.32, -0.25]	[-0.05, -0.01]	[0.09, 0.13]	[0.16, 0.24]

Notes: Numbers are for households with 2 members. For each choice occasion we compute the market-share weighted mean own price elasticity. Numbers shows average of this own price elasticity. We measure expenditure as the households' mean weekly grocery expenditure over the year. 95% confidence intervals are shown in brackets.

Table 4.7: *Compensating variation from specific tax for household size 2: Cubic and linearised utility*

Specification	Mean compensating variation	Average deviation from mean compensating variation for quartile of grocery expenditure distribution:			
		1	2	3	4
Cubic utility	2.49 [2.34, 2.69]	-0.70 [-0.80, -0.64]	-0.06 [-0.13, -0.03]	0.28 [0.23, 0.35]	0.49 [0.39, 0.61]
Linearised utility	2.49 [2.31, 2.68]	-0.71 [-0.80, -0.62]	-0.07 [-0.14, -0.03]	0.28 [0.24, 0.34]	0.50 [0.40, 0.62]

Notes: Numbers are for households with 2 members. Numbers give compensating variation for the average household associated with the simulated specific tax. We measure expenditure as the households' mean weekly grocery expenditure over the year. Numbers are for a calendar year. 95% confidence intervals are shown in brackets.

4.8.2 Comparison of specific and ad valorem taxes

In Table 4.8 we summarise the aggregate annual effects of both the specific tax and the ad valorem tax. The results are computed using our baseline cubic utility specification. Differences in the two tax's aggregate impacts are driven by the fact that i) the specific tax is a function of product saturated fat content while the ad valorem tax is a function both product saturated fat content and price; and ii) equilibrium pass-through of the taxes differs. On average pass-through of the specific tax is 102% while pass-through of the ad valorem tax is 54%. The ad valorem tax makes increasing producer prices more costly (in terms of lost demand) than a specific tax does because firms must increase consumer prices by more under the ad valorem tax to achieve a given increase in producer prices. Because this affects all firms and because of firms' strategic interactions, this results in lower equilibrium prices. Therefore, while we set the specific and ad valorem tax rates to achieve the same reduction (of 20%) in saturated fat purchases when producer prices are held fixed, the pricing response of firms results in different effects on aggregate saturated fat purchases in equilibrium.

The top panel of Table 4.8 summarises the impact of the taxes on consumer expenditure on butter and margarine, firm profits, consumer compensating variation and welfare. Here compensating variation takes account of preferences consumers have at the point of purchase – if, for instance, consumers suffer from internalities any con-

sumer welfare loss measured here could ultimately be overturned. Similarly, welfare accounts for changes in firm profits, consumer welfare and tax revenue raised. It does not account for savings in any externalities averted due to diminished saturated fat consumption.

Consumer spending in the market is around £530m. The specific tax lowers this by 8.9% while the ad valorem tax lowers it by 8.3%. Both taxes lead to a reduction in firm profitability, by 15.1% in the case of the specific tax and 16.3% in the case of the ad valorem tax. While the specific tax leads to a marginally lower reduction in profits than the ad valorem tax, it leads to a much larger reduction in consumer welfare. Compensating variation is £56m for the specific tax against £33m for the ad valorem tax. As both taxes raise similar amounts of tax revenue (£45-46m) the welfare loss associated with the specific tax is greater than that associated with the ad valorem tax.

The second panel of Table 4.8 summarises the impact of the taxes on consumer saturated fat purchases. Collectively consumers purchase around 43m kg of saturated fat in the form of butter and margarine annually. The specific tax reduces this by 21.5% and the ad valorem tax reduces it by 14.2%. The differences are driven by the differential pass-through of the two forms of tax.

In the final panel we compare the average consumer and welfare cost per 1kg reduction in saturated fat. The consumer cost is based on compensating variation and the welfare cost is based on compensating variation plus fall in firm profits minus tax revenue raised. These numbers give an idea of the cost effectiveness of the two forms of tax in lowering consumer saturated fat consumption. In both cases the cost of reducing fat is around 50 pence higher per kg for the specific tax. Therefore, while the specific tax is more effective at lowering saturated fat purchases (due to higher pass-through to consumer prices), it leads to sufficiently higher costs to consumers to result in a marginally higher average welfare cost per unit of saturated fat reduction than the comparable ad valorem tax.

Table 4.8: *Aggregate effect of tax*

	Pre tax	Specific tax	Ad valorem tax
Expenditure (£m)	530.01	482.61	486.05
(% change)		-8.94	-8.29
Profits (£m)	241.83	205.17	202.48
(% change)		-15.16	-16.27
Compensating variation (£m)		55.51	33.36
Tax revenue (£m)		46.26	45.44
Welfare (£m)		-45.91	-27.27
Saturated fat (m kg)	42.65	33.50	36.57
(% change)		-21.45	-14.24
Consumer cost per 1kg fat		6.07	5.49
Welfare cost per 1kg fat		5.02	4.49

Notes: Results are grossed up to British population and are for one year. Welfare refers to tax revenue minus the sum of compensating variation and fall in firm profits.

5 Conclusion

In this paper we study the empirical importance of relaxing functional form restrictions commonly placed on how income or total expenditure enter logit demand models. We focus on a market for a small budget share product category – butter and margarine – and show that even here allowing for flexibility in the relationship between grocery expenditure and demands is important. For small budget share goods it is standard to assume the marginal utility of total expenditure (or income) is constant, allowing for expenditure to enter the model only in an ad hoc way as a linear preference shifter. We show in our setting that this approach does a good job of recovering market level average elasticities, marginal costs, pass-through and consumer welfare. This is encouraging, as if researchers only have access to market level rather than consumer level data, incorporating more flexibility in the relationship may be difficult. However, compared with a more general model in which total expenditure enters utility in a flexible theoretically coherent fashion as a polynomial, the standard model fails in recovering distributional aspects of demand and welfare effects.

We also provide evidence on the impacts of a “fat tax” levied on products’ saturated fat contents. The flexibility we introduce in consumer utility is crucial for studying the

equilibrium pricing response to tax. Logit models that assume utility is linear in total expenditure (or income) impose that consumer level demands are log-concave. This functional form assumption can place restrictions on equilibrium tax pass-through. Our polynomial utility model relaxes these restrictions and therefore we can be more confident our results are not prejudiced by functional form assumption. We find that pass-through of a specific tax is greater than an ad valorem tax and that therefore the former is more effective at leading to reductions in saturated fat. However, this comes at the cost of higher reductions in consumer welfare.

In our application the marginal utility of grocery expenditure is clearly non-constant. However, because we consider a small market share good, the change in price induced by the tax is small relative to $y - p_j$. The policy change itself induces a small income effect. This is important in understanding why the preference shifter model successfully recovers the aggregate consumer welfare change. Similarly, because we find that the curvature of household level demands under the polynomial utility model is similar to that in the more restrictive models in which utility is linear in price, the preference shifter model is able to recover the same pattern of pass-through as the more general model. In applications in which a tax induces a price change that is large relative to $y - p_j$, or in which the curvature of individual demands is less well captured by the log-concave shape of a logit model with utility linear in price, the preference shifter model would do less well at replicating the average results of the polynomial utility specification.

In applications to product categories comprising large shares of consumers' budgets, flexibly modelling how total expenditure enters utility is likely to be even more important than in our application. In such markets, price changes are more likely to be large enough to induce significant income effects. In applications involving large budget share items (e.g. cars) it has been common to allow for income effects through use of the log utility formulation. Our results suggest this specification may be overly restrictive and insufficiently flexible to capture the true variation in the marginal utility and should be tested against more flexible alternative specifications.

In this study we use rich data and exploit both nonlinear pricing within brands and cross-sectional variation in mean (over the year) weekly grocery expenditure to pin down the relationship between demand for butter and margarine and both price and expenditure. We control for many important potentially confounding factors including

advertising expenditure, brand fixed effects, and household size. Conditional on these covariates, we rely on the assumptions that both prices and expenditure are exogenous.

There are several promising avenues for future work. In many cases, researchers do not have access to rich data such as ours. For example, in many applications, data on expenditure are not available. Our results suggest that in such cases, researchers should consider specifications with nonlinear random coefficients that capture the distribution of unobserved expenditure. In other applications, the assumption of exogenous total expenditure or income may be less plausible than in our case. The use of panel data to isolate within consumer time series variation could potentially enable researchers to identify income effects. Price endogeneity may also be an important concern in other settings; research is needed into methods to deal with this concern while allowing for nonlinear conditional utility.

References

- Anderson, S. P., A. De Palma, and B. Kreider (2001). Tax incidence in differentiated product oligopoly. *Journal of Public Economics* 81(2), 173–192.
- Bajari, P. and C. L. Benkard (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of political economy* 113(6), 1239–1276.
- Banks, J., R. Blundell, and A. Lewbel (1997, November). Quadratic engel curves and consumer demand. *Review of Economics and Statistics* 79(4), 527–539.
- Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 63(4), 841–890.
- Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy* 112(1), 68–105.
- Besley, T. and H. Rosen (1999). Sales taxes and prices: an empirical analysis. *National tax journal* 52(2), 157–178.
- Bhattacharya, D. (2015). Nonparametric welfare analysis for discrete choice. *Econometrica* 83(2), 617–649.
- Browning, M. and C. Meghir (1991). The effects of male and female labor supply on commodity demands. *Econometrica*, 925–951.
- Dagsvik, J. K. and A. Karlström (2005). Compensating variation and hicksian choice probabilities in random utility models that are nonlinear in income. *The Review of Economic Studies* 72(1), 57–76.
- Deaton, A. and J. Muellbauer (1980). An Almost Ideal Demand System. *The American Economic Review* 70(3), 312–326.
- Delipalla, S. and M. Keen (1992). The comparison between ad valorem and specific taxation under imperfect competition. *Journal of Public Economics* 49(3), 351–367.
- Delipalla, S. and O. O’Donnell (2001). Estimating tax incidence, market power and market conduct: the european cigarette industry. *International Journal of Industrial Organization* 19(6), 885–908.

- Draganska, M., M. Mazzeo, and K. Seim (2009). Beyond plain vanilla: Modeling joint product assortment and pricing decisions. *Quantitative Marketing Economics* 7(2), 105–146.
- Dubé, J., G. Hitsch, and P. Manchanda (2005). An empirical model of advertising dynamics. *Quantitative Marketing and Economics* 3, 107–144.
- Fabinger, M. and E. G. Weyl (2015). A tractable approach to pass-through patterns. *mimeo*.
- Fabra, N. and M. Reguant (2014). Pass-through of emissions costs in electricity markets. *The American Economic Review*, Forthcoming.
- Goldberg, P. K. and R. Hellerstein (2008). A structural approach to explaining incomplete exchange-rate pass-through and pricing-to-market. *The American Economic Review*, 423–429.
- Goldberg, P. K. and R. Hellerstein (2013). A structural approach to identifying the sources of local currency price stability. *The Review of Economic Studies* 80(1), 175–210.
- Goldberg, P. K. and F. Verboven (2001). The evolution of price dispersion in the european car market. *The Review of Economic Studies* 68(4), 811–848.
- Hagenaars, A. J., K. De Vos, M. Asghar Zaidi, et al. (1994). Poverty statistics in the late 1980s: Research based on micro-data.
- Hamilton, S. F. (2009). Excise taxes with multiproduct transactions. *American Economic Review* 99(1), 458–471.
- Hausman, J. and W. Newey (2014). Individual heterogeneity and average welfare. *Cemmap working paper*.
- Herriges, J. A. and C. L. Kling (1999). Nonlinear income effects in random utility models. *Review of Economics and Statistics* 81(1), 62–72.
- Kenkel, D. S. (2005). Are alcohol tax hikes fully passed through to prices? evidence from alaska. *American Economic Review*, 273–277.
- Kim, D. and R. W. Cotterill (2008). Cost pass-through in differentiated product markets: The case of us processed cheese. *The Journal of Industrial Economics* 56(1), 32–48.
- Lewbel, A. and K. Pendakur (2009). Tricks with Hicks: The EASI Demand System. *The American Economic Review* 99(3), pp. 827–863.

- McFadden, D. (1981). Econometric models of probabilistic choice among products. In C. F. Manski and D. McFadden (Eds.), *Structural analysis of discrete data with econometric applications*, pp. 198–272. Cambridge: MIT Press.
- McFadden, D. (1999). Computing willingness-to-pay in random utility models. *Trade, Theory, and Econometrics: Essays in Honour of John S. Chipman* 15, 253.
- McFadden, D. and K. Train (2000). Mixed mnl models for discrete response. *Journal of applied Econometrics* 15(5), 447–470.
- Nakamura, E. and D. Zerom (2010). Accounting for incomplete pass-through. *The Review of Economic Studies* 77(3), 1192–1230.
- Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, 395–421.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69(2), 307–342.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* 110(4), 705–729.
- Seade, J. (1985). Profitable cost increases and the shifting of taxation: Equilibrium responses of markets in oligopoly. *Warwick economic research papers*.
- Small, K. A. and H. S. Rosen (1981). Applied welfare economics with discrete choice models. *Econometrica*, 105–130.
- Villas-Boas, S. B. (2007). Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies* 74(2), 625–652.
- Weyl, E. G. and M. Fabinger (2013). Pass-through as an economic tool: principles of incidence under imperfect competition. *Journal of Political Economy* 121(3), 528–583.