

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

Rachel Griffith
Lars Nesheim
Martin O'Connell

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP23/15

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

Rachel Griffith, Lars Nesheim and Martin O’Connell*

June 4, 2015

Abstract

Random utility models are widely used to study consumer choice. The vast majority of applications make strong assumptions about the marginal utility of income, which restricts income effects, demand curvature and pass-through. We show that flexibly modeling income effects can be important, particularly if one is interested in the distributional effects of a policy change, even in a market in which, a priori, the expectation is that income effects will play a limited role. We allow for much more flexible forms of income effects than is common and we illustrate the implications by simulating the introduction of an excise tax.

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

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 and University College London, martin_o@ifs.org.uk

1 Introduction

Random utility models are widely used to study consumer choice among differentiated products. When using such models, it is common to make strong assumptions about the marginal utility of income. Such assumptions help with model tractability, simplify analysis of counterfactual equilibria and simplify welfare calculations. It is well understood that these assumptions are restrictive. They place strong restrictions on income effects, on the curvature of demand, and hence on predictions of pass-through (see, *inter alia*, McFadden (1999), Herriges and Kling (1999), Weyl and Fabinger (2013) and Fabinger and Weyl (2014)). Nevertheless, it is commonly believed that for small budget share product categories, the assumption of a constant marginal utility of income is a reasonable approximation. Many such applications allow income to enter in an *ad hoc* form as a “preference shifter”. Alternatively, for larger budget share products, it is common to include income or total expenditure in log form.

In this paper, we show that flexibly modeling income effects can be important, particularly if one is interested in the distributional effects of a policy change, even in a market in which, *a priori*, the expectation is that income effects will play a limited role. We allow for much more flexible forms of income effects than is common in the applied discrete choice demand literature and, to highlight the implications of flexibly incorporating income effects, 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 standard to the literature.

The existing literature that uses logit models to estimate consumer demand for differentiated goods has typically made one of two assumptions regarding the nature of income effects. Most commonly among papers focusing on product categories that comprise a small budget share, researchers have assumed that utility is linear in income (or expenditure) minus price, conditional on selecting a given option (see, Nevo (2001), Villas-Boas (2007)). Under this assumption, income 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 income and purchase patterns, researchers often include income in a reduced form way as a preference shifter, which linearly shifts the coefficient on price. The second common assumption made in the literature, usually when the application relates to product categories that comprise a large share of consumers' budgets, is that conditional utility is linear in the log of income (or expenditure) minus price (for example, Berry et al. (1995), Goldberg and Verboven (2001) and Petrin (2002)). This specification incorporates income effects into the model, but does so in a restrictive way. At the consumer level, the conditional marginal utility of income is inversely proportional to income.

We show that neither of these standard models can fully replicate the results we obtain with our more general model with flexible income effects. The log utility model produces estimates of demand, pass-through and welfare that are biased, both at the average and across the total expenditure distribution. Although the log utility model does allow for income effects, the restrictive function it imposes on the conditional marginal utility of income is strongly rejected in more flexible specifications. In contrast, the preference shifter model yields estimates of market level average quantities, such as tax pass-through, average price elasticities and aggregate welfare effects, that are similar to a model with flexible income effects. However, it fails to fully recover variation in price sensitivity and welfare effects across the expenditure distribution.

The preference shifter model does not admit income effects. However, it does allow for some cross sectional correlation in demand patterns and welfare effects with total expenditure through including the latter as a linear shifter of the price coefficient. We show that when income effects exist in demand, but when the counterfactual of interest involves relatively small price changes that do not themselves induce large income effects, a simple modification to the standard preference shifter model, which involves interacting price with higher order expenditure terms, can do a very good job of replicating the distributional results found with the full model with flexible income effects. 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 linearized model - which can be approximated by interacted price with functions of total expenditure (or income) - performs well when analyzing impacts of small price changes.

On the other hand, if one is interested in studying a policy reform that shifts prices by a large amount relative to total expenditure (for instance in a market for large budget share goods) or if one is interested in understanding how consumers would respond to a policy that changed total expenditure or income, it is necessary to estimate the model with flexible income effects to obtain unbiased estimates of the effect of the reform.

We investigate the empirical relevance of these issues in a market where income effects would not seem to be a major concern, the butter and margarine market. In the UK, this market represents only about 1% of average total grocery expenditure. Nonetheless, in common with many other product categories, consumer purchase patterns are strongly related to total grocery expenditure. In particular, higher total grocery expenditure is associated with a higher probability of purchase, and, conditional on purchase, selecting a relatively high priced option. We estimate demand in this market allowing a cubic spline to capture the structural relationship between a consumer's net expenditure and their utility from purchasing an option, and we show that the relationship is nonlinear and approximately cubic. Standard specifications are unable to recover the distributional consequences of introducing a tax, but we show the correctly linearized model is successful in doing this..

The fact that, in the case of small price changes, a linear approximation of the utility model with flexible income effects 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 income effects can be considerably more costly. In particular, in the model with income effects 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) or Dagsvik and

Karlström (2005). Recently Bhattacharya (2015) has shown how to estimate the marginal distribution of compensating variation non-parametrically when interest centers on the impact of a change in the price of a single good. However, in differentiated product markets in which interest typically centers 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 nature of income effects 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.

Assumptions about income effects 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 stylized 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 form of income effects 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. We explore the importance of relaxing demand curvature restrictions through allowing for flexible income effects when assessing equilibrium pass-through of a tax to consumer prices.

Our work is related to a large literature that estimates pass-through of cost shocks and taxes to prices. A series 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 excise 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. Another feature of this literature has been to highlight that nominal rigidities may be important in generating delayed adjustment to shocks, although they are less important in determining long-run pass-through. 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 modeling income effects in random utility models and their implication 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 results from an empirical example. We compare how different forms of income effects in demand impact on the consumer welfare effects and pass-through of an excise tax. A final section concludes.

2 Consumer level demand

We consider a random utility model of consumer choice (see McFadden (1981)). The consumer has a total 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. For instance, in applications to a particular grocery product category, y may be total grocery expenditure. The consumer makes a discrete choice about which alternative $j \in \{0, 1, \dots, J\}$ to purchase and spends their remaining budget on other groceries. We denote the price of option j as p_j . Option $j = 0$ denotes the 'outside option' and $p_0=0$. Option j has associated with it a vector of observable product characteristics \mathbf{x}_j and unobservable characteristics 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 . 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.

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; it is non-increasing in prices, non-decreasing in total budget, homogeneous of degree zero in all prices and total budget, quasi-convex in prices and continuous in prices and total budget. 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 $\boldsymbol{\varepsilon}$ are 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 maximization 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 is additive and is drawn from a generalized 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:

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

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

The probability the consumer selects option j ,

$$P_j = \Pr [U(y - p_j, \mathbf{x}_j, \varepsilon_j) \geq U(y - p_k, \mathbf{x}_k, \varepsilon_k) \ \forall k],$$

under (2.2) 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 restricts this specification even further by imposing that the marginal utility of income is constant:

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

This specification rules out income effects. At the consumer level, when comparisons are made across options, y differences out of the model. 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 total budget (and possibly also with other demographic variables):

$$\alpha = \alpha_0 + \alpha_1 y + \nu \quad (2.5)$$

where ν is a random coefficient. This “preference shifter” model has no 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 expenditure patterns do vary cross-sectionally with income or total expenditure.

Papers that do allow for income effects 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

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

In this case the conditional marginal utility of income is given by $\frac{\alpha}{y - p_j}$; the conditional marginal utility of income is inversely proportional to income.

In the following sections, we explore the importance of allowing for richer forms of income effects. We first discuss implications for consumer welfare and for the curvature of consumer demand. Then we develop an empirical application to the market for butter and margarine and show that income effects are important for estimating accurately how individual demand elasticities and welfare effects depend on y .

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 organization, 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 income in the model as a preference shifter (as in equation (2.5)).¹ 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, ϵ , are realized, and can be computed (conditional on realizations 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 realized (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 ϵ . In this case, aggregate welfare is the sum of the individual expected utilities. Conversely, if there is no uncertainty for the consumer over ϵ but rather the logit shocks simply represent cross-sectional unobserved heterogeneity, then consumer welfare changes should be based on an ex post criterion based on the individual consumer's *realized utility*. In this case, aggregate welfare is the sum or average of the individual's realized

¹Petrin (2002) is an exception. He uses the log specification given by equation (2.6), and he estimates the consumer welfare effects of the introduction of minivans to the automobile market.

utilities. We present results adopting the latter perspective, based on realized 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.² Individual level compensating variation, cv , associated with the price change satisfies

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

Individual cv depends on $\boldsymbol{\varepsilon}$ 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 characterize 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 eliminate simulation error in computing CV at the cost of much higher computational effort.

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 (see, *inter alia*, Seade (1985), Anderson et al. (2001) and Weyl and Fabinger (2013)). Weyl and Fabinger (2013) emphasize that, in

²The 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.

the context of a monopolist or symmetrically differentiated single product firm oligopoly, the curvature of the log of demand is key. A simple example illustrates the point.

Consider a single product monopolist with constant marginal cost, c . Let the demand curve be $q(p)$. Optimization 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 ($\frac{d^2 \ln q}{dp^2} < 0$). 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 may 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 bearing on the permissible curvature of the log of market demand, and therefore on pass-through.

We develop these ideas in the context of the market demand curve allowing for individual heterogeneity in both income and preferences. Let each consumer be indexed by (y, θ) where as discussed above y measures income and θ measures all other observable and unobservable consumer attributes that enter into utility. Normalizing the size of the market to be one, the market demand curve for option j is then given by:

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

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:

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

The curvature of the log of market demand depends on two terms. The first term is the probability weighted average of the second derivatives of log individual demand. The second term is the probability 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 of income 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 Illustrative application

To illustrate the potential importance of modeling income effects in a flexible way we provide an example using the UK market for butter and margarine. We have purposely chosen a market that represents a small share of expenditure (this market accounts for just over 1% of households' regular grocery expenditure). The expectation is that income effects play a limited role in this market. We estimate demand under a number of different assumptions about the nature of income effects. We compute individual and market level demand elasticities and simulate the impact of an excise tax. We compare the tax pass-through and consumer welfare predictions of the various specifications.

4.1 Consumers

Let i index consumers and t denote time. The index $j \in \{1, \dots, J\}$ indexes butter and margarine products. We define a product as a brand-pack size combination and index brands by $b = 1, \dots, B$. Product $j = 0$ is the outside option. The product characteristics are given by $\mathbf{x}_{jt} = (a_{bt}, z_j, \mathbf{w}_b, \xi_b)$. a_{bt} measures advertising expenditure for brand b in period t . z_j is pack size, \mathbf{w}_b is a vector of observable brand characteristics and ξ_b is an unobserved (by the econometrician) brand characteristic. In our application, a_{bt} varies over time but not within brand, z_j varies within brand but not over time, and (w_b, ξ_b) do not vary over time. We discuss how this variation contributes to identification in Section 4.2 below.

We assume preferences for groceries are weakly separable from other goods and measure y_i as consumer i 's average weekly grocery expenditure over a calendar year. By 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).

We assume utility from selecting butter or margarine product j takes the form

$$U_{ijt} = f(y_i - p_{jt}; \alpha_i) + \gamma a_{bt} + \lambda_i z_j + \beta'_i \mathbf{w}_b + \xi_b + \varepsilon_{ijt}, \quad (4.1)$$

and utility from selecting the outside option is given by

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

We specify a number of different forms for $f(y_i - p_{jt}; \alpha_i)$.

- Polynomial utility

$$f(y_i - p_{jt}; \alpha_i) = \alpha_i^{(1)}(y_i - p_{jt}) + \sum_{n=2}^N \alpha^{(n)}(y_i - p_{jt})^n$$

- Linear utility

$$f(y_i - p_{jt}; \alpha_i) = \alpha_i^{(1)}(y_i - p_{jt})$$

- Preference shifter

$$f(y_i - p_{jt}; \alpha_i) = (\alpha_i^{(1)} + \alpha^y y_i) p_{jt}$$

- Linearized cubic utility

$$f(y_i - p_{jt}; \alpha) = -(\alpha_i^{(1)} + \alpha^{y1} y_i + \alpha^{y2} y_i^2) p_{jt}$$

- Log utility

$$f(y_i - p_{jt}; \alpha_i) = \alpha_i^{(1)} \ln(y_i - p_{jt})$$

- Spline utility

$$f(y_i - p_{jt}; \alpha_i) = \sum_{k=1}^K a_i^{(k)} B^{(k)}(y_i - p_{jt})$$

where $B^{(k)}(y_i - p_{jt})$ are a set of cubic B-splines with $K - 2$ knots placed at the extremes and at the equally spaced percentiles of the expenditure distribution.

The first and last specifications include $y_i - p_{jt}$ in utility in a flexible nonlinear way, and therefore admit flexible forms of income effects. We focus our analysis on the polynomial utility case to highlight the role of income effects and to compare more easily with specifications commonly used in the literature. We find empirically that the estimated cubic polynomial utility model closely mimics the estimated spline utility model. In our application, the spline utility estimates are similar to the cubic utility estimates but are more costly to compute.

The linear utility specification rules out income effects; it assumes that the marginal utility of income is constant. It also does not allow purchase patterns to be correlated with total expenditure. The preference shifter specification also assumes the marginal utility of income is constant but does allow the price parameter to shift linearly with total expenditure across households. The log utility model specifies utility to be nonlinear in $y_i - p_{jt}$ and so admits income effects, but $y_i - p_{jt}$ enters utility much less flexibly than in the polynomial or

spline utility specifications. We show below that this specification performs very poorly in our application.

The coefficients on the first order $y_i - p_{jt}$ term, pack size and observable brand attributes are allowed to vary across consumers. In particular, we model these coefficients as:

$$\begin{aligned}\alpha_i^{(1)} &= \alpha_0^{(1)} + \nu_i^{(\alpha)} \\ \lambda_i &= \lambda_0 + \lambda_1 d_i \\ \beta_i &= \beta_0 + \beta_1 d_i + \nu_i^{(\beta)}\end{aligned}$$

where d_i represents observable consumer demographics. We assume $\nu = (\nu_i^{(\alpha)}, \nu_i^{(\beta)})' \sim N(\mathbf{0}, \Sigma)$ and is uncorrelated with y_i and d_i . We control for the unobserved brand attribute by including brand fixed effects. β_0 is therefore absorbed into the brand fixed effects. Unobserved preference heterogeneity incorporated through the random coefficients allows for correlations in the unobserved portion of consumer utility across options and choice occasions and is crucial in enabling logit choice models to capture realistic substitution patterns across options (see, for instance, Train (2003) and Berry et al. (1995)).

4.2 Identification

A common concern in empirical demand analysis is that the ceteris paribus impact of price on demand may not be identified because there may be unmeasured demand shocks that are correlated with price. The most common concern is that price might be correlated with an unobserved product effect, either some innate unobserved characteristic of the product or some market specific shock to demand for the product. Failure to control for the unobserved product effect will lead to inconsistent estimates of the true price effect.

Due to features of the UK market for butter and margarine and the richness of product level data, we believe the identification strategy proposed in Bajari and Benkard (2005) is reasonable in this setting. In particular, we exploit two forms of price variation to identify

the marginal utility of income. Firstly, conditional on brand fixed effects and advertising, we argue that it is reasonable to assume that there is no variation in ξ_b (i.e. the relative desirability of one brand over a second does not fluctuate throughout the year). We also argue that the fact that the UK retail food market is characterized by close to national pricing limits the risk that the individual specific demand innovations, ε_{ijt} , conditional on advertising, are correlated with prices.³ We therefore exploit time series variation in the set of prices. Secondly, we exploit differential nonlinear pricing within brands. This second source of price variation is important in our case and is not exploited in most applications in the literature. Unlike most applications our very detailed data allows us to model demand at the product, rather than brand, level. We are therefore able to include a full set of brand and pack size effects, while exploiting cross-sectional price variation due to differential nonlinear pricing within brands.

In the specifications in which utility is a nonlinear function of net expenditure, $y_i - p_{jt}$, we are able to exploit an additional source of variation, the large cross-sectional variation in grocery expenditure across households, to identify the marginal utility of income. This large cross-sectional variation allows us to estimate a very flexible model of income effects. To avoid endogeneity concerns about trip-level grocery expenditure, we measure household expenditure as the household's average weekly expenditure over the course of a calendar 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 errors in butter and margarine demand (a demand shock leading to the purchase of a particularly expensive butter product would be correlated, all else equal, with higher trip grocery expenditure). A second issue might be that much of the high frequency variation in trip level expenditure might reflect planning decisions related to how many shopping trips to undertake in a given period of time and would not be informative of income effects. Use of average weekly expenditure minimizes these concerns and ensures that we only exploit variation in total expenditures

³In the UK most supermarkets implement a national pricing policy following the Competition Commission's investigation into supermarket behavior (Competition Commission (2000)).

which reflects long run expenditure decisions. This will be valid if unobserved preferences that affect substitution within the butter and margarine market are independent from factors that affect average weekly expenditure on all groceries.

4.3 Firm competition

Let $f = \{1, \dots, F\}$ index firms and F_f denote the set of products owned by firm f . We assume that firms 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)).

Normalizing the size of the market to be one, 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.

4.4 Counterfactual

We simulate the introduction of an excise 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^e satisfies:

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

In an appendix we show that our analysis and results yield similar conclusions if instead we consider an ad valorem tax (τ_{av}), such that a counterfactual equilibrium price vector \mathbf{p}_t^e satisfies:

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

4.5 Data

We apply the model to the UK market for butter and margarine. We use purchase data on 10,012 households from Kantar WorldPanel for calendar year 2010. For each household we observe all grocery products that are bought and taken into the home. We define a ‘choice occasion’ as a household’s weekly grocery purchases. We use information on five randomly chosen choice occasions for each household (50,060 in total) to estimate the model. Households purchase a butter or margarine product on 34% of choice occasions. On these choice occasions households on average spend around £1.35 (or 3.5% of their grocery expenditure) on butter and margarine. On the remaining occasions they select the outside option of not purchasing butter or margarine.

We measure grocery expenditure, y_i , as a household’s mean weekly grocery expenditure over 2010. As discussed in Section 4.2 using average weekly expenditure ensures that our expenditure measure is not correlated with idiosyncratic trip-specific demand shocks. Average expenditure in the sample is £40. The 10th percentile of the distribution is £20 and the 90th percentile is £63.

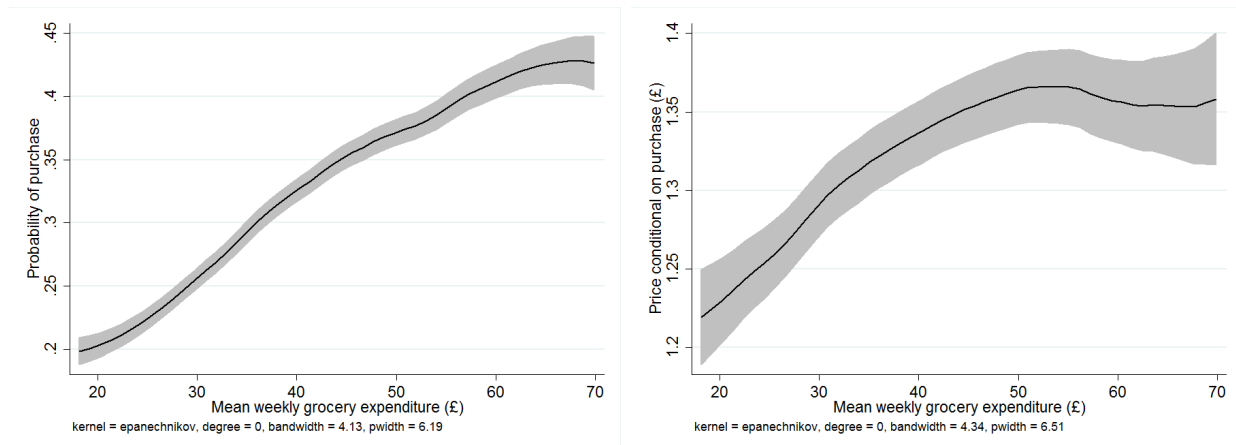
We provide some reduced form evidence that, despite the fact that butter and margarine are small budget share items, household butter and margarine purchase behavior is correlated

with their mean weekly grocery expenditure. The left side panel of Figure 1 shows results from a non-parametric regression of the average probability that a household purchases any butter or margarine product on a choice occasion on its mean weekly grocery expenditure. The right side panel displays results from a non-parametric regression of the average price a household pays for butter or margarine, conditional on purchasing, on its mean weekly grocery expenditure. The figures show 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. Table 1 shows that this pattern remains after conditioning on household size. The pattern is not simply a reflection of correlations between household size, expenditure, and purchase patterns.

Figure 1: *Correlation of purchase patterns with mean weekly grocery expenditure:*

A) *probability of purchase*

B) *price paid conditional on purchase*



Notes: The figures display results from weighted kernel regressions across 10,012 households. The weights 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.

Table 1: *Variation in purchase behavior with mean weekly grocery expenditure by household size*

| Quartile of expenditure distribution | Probability of purchasing butter margarine | Price paid conditional on purchase |
|--------------------------------------|--|------------------------------------|
| <i>One person households</i> | | |
| 1 | 0.19 | 0.22 |
| 2 | 0.23 | 0.29 |
| 3 | 0.23 | 0.32 |
| 4 | 0.30 | 0.43 |
| <i>Two person households</i> | | |
| 1 | 0.26 | 0.32 |
| 2 | 0.35 | 0.46 |
| 3 | 0.39 | 0.56 |
| 4 | 0.42 | 0.60 |
| <i>Three person households</i> | | |
| 1 | 0.27 | 0.33 |
| 2 | 0.35 | 0.48 |
| 3 | 0.38 | 0.52 |
| 4 | 0.45 | 0.65 |
| <i>Four person households</i> | | |
| 1 | 0.28 | 0.34 |
| 2 | 0.36 | 0.47 |
| 3 | 0.42 | 0.54 |
| 4 | 0.45 | 0.63 |

Notes: Quartiles are defined for the distribution of mean weekly grocery expenditure within each household size category.

The butter and margarine market in the UK has an oligopolistic structure. There are eight main firms in the market. Unilever is the largest, 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 first four columns of Table 2 list the firms 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. In most cases, a product (i.e. an option in a consumer's choice set) is defined as a brand-pack size combination.⁴ Column five shows the quantity share of each product.

⁴In a few instances a brand-pack size contains two products - a salted and unsalted version.

Column six shows the mean market price computed as the transaction weighted mean price in each month. The remaining columns show how the product characteristics we include in the model vary across products. Characteristics include whether the product is butter or margarine, the amount of saturated fat per 100g and monthly advertising expenditure for the brand (in addition to pack size and brand effects).

Table 2: *Products and characteristics*

| 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.

4.6 Estimates

We estimate the five specifications - polynomial utility, linear utility, preference shifter, log utility, and spline utility - outlined in Section 4.1 using maximum likelihood.⁵ In Table 3 we report the coefficient estimates.⁶ For the polynomial utility specification, we specify utility as a third order polynomial of net expenditure, $y - p_j$. As shown in Figure 2 this provides sufficient flexibility to capture the shape of the conditional marginal utility of income implied by the estimates of the spline utility specification.

The top panel presents estimates of the random coefficients. For each specification we model the coefficient on the first order net expenditure term as a random coefficient, meaning we allow 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, but we allow the mean to shift with whether the main shopper has a body mass index indicating he or she is obese. We assume that the random coefficients are joint normally distributed and allow for correlation between the coefficients. Direct interpretation of the $y - p_j$ coefficients is difficult. We simply note that the means of all $y - p_j$ coefficients are statistically significant, as are all the higher order, interaction, variance and covariance parameters.

The bottom section of the table shows the coefficient estimates for the non-random coefficients. In each case the advertising coefficient is statistically insignificant indicating little evidence that butter and margarine advertising has a strong contemporaneous impact on demand. We interact pack size effects with household size, which captures the fact that larger households are more likely to select large pack sizes. We also interact pack size with a dummy indicating whether the main shopper is obese. Three of the four specifications indicate obese main shoppers have a statistically significant preference for 500g and 1kg pack

⁵We use Gauss-Hermite quadrature rules to eliminate simulation error when computing the likelihood function.

⁶For brevity, we do not show spline results in Table 3. Results from the spline utility model are statistically indistinguishable from the cubic utility model. In Figure 2 we illustrate this point by graphing the estimated conditional marginal utility of income from both the spline utility and cubic utility models.

sizes over the smaller 250g pack size (the coefficients are not statistically significant in the log utility model). Like the butter dummy, a product's saturated fat content per 100g is collinear with the brand effects. Therefore we include this attribute interacted with the obese dummy. In all four specifications the obese-butter interaction is positive and statistically significant and the obese-saturated fat interaction is not statistically significant, indicating obese consumers have a stronger preference than other consumers for butter, but conditional on their preference for butter, they do not prefer products with higher saturated fat.

Table 3: *Coefficient estimates*

| | Polynomial utility | | Linear utility | | Preference shifter | | Log utility | |
|----------------------------------|----------------------|----------------|----------------------|----------------|----------------------|----------------|----------------------|----------------|
| | Coefficient estimate | Standard error | Coefficient estimate | Standard error | Coefficient estimate | Standard error | Coefficient estimate | Standard error |
| Random coefficients | | | | | | | | |
| <i>Mean terms</i> | | | | | | | | |
| (y-p) | 3.8134 | 0.0888 | 2.5238 | 0.0507 | 3.1991 | 0.0602 | 4.4524 | 0.1443 |
| <i>Higher order terms</i> | | | | | | | | |
| (y-p) ² | -0.2352 | 0.0164 | | | | | | |
| (y-p) ³ | 0.0115 | 0.0012 | | | | | | |
| <i>Interaction terms</i> | | | | | | | | |
| (y-p)*y | | | | | -0.1512 | 0.0067 | | |
| Butter*Obese | 0.1301 | 0.0713 | 0.1482 | 0.0720 | 0.1406 | 0.0712 | 0.1482 | 0.0975 |
| <i>Variance-covariance terms</i> | | | | | | | | |
| Var(y-p) | 0.7410 | 0.0149 | 0.7851 | 0.0151 | 0.7515 | 0.0148 | 6.0865 | 0.3788 |
| Var(ln(y-p)) | | | | | | | 4.5884 | 0.2307 |
| Var(Butter) | 1.7198 | 0.0333 | 1.8075 | 0.0324 | 1.7332 | 0.0331 | | |
| Cov(y-p, Butter) | 1.4204 | 0.0466 | 1.3610 | 0.0448 | 1.4083 | 0.0461 | | |
| Cov(ln(y-p), Butter) | | | | | | | 2.9557 | 0.2026 |
| Fixed coefficients | | | | | | | | |
| Advertising | -0.0307 | 0.0224 | -0.0305 | 0.0224 | -0.0306 | 0.0224 | -0.0180 | 0.0321 |
| 500g | 3.0214 | 0.0640 | 2.7368 | 0.0634 | 3.0053 | 0.0640 | 2.0281 | 0.0692 |
| 1kg | 3.6145 | 0.1190 | 3.1601 | 0.1178 | 3.6073 | 0.1191 | 1.2973 | 0.1187 |
| 500g*HHsize | 0.0796 | 0.0109 | 0.1887 | 0.0102 | 0.0840 | 0.0109 | 0.0109 | 0.0134 |
| 1kg*HHsize | 0.2629 | 0.0221 | 0.4343 | 0.0211 | 0.2659 | 0.0221 | 0.1974 | 0.0266 |
| 500g*Obese | 0.1321 | 0.0298 | 0.1730 | 0.0303 | 0.1409 | 0.0299 | 0.0307 | 0.0381 |
| 1kg*Obese | 0.1630 | 0.0592 | 0.2255 | 0.0598 | 0.1760 | 0.0596 | 0.0314 | 0.0740 |
| Satfat*Obese | 0.0397 | 0.1178 | 0.0582 | 0.1181 | 0.0465 | 0.1178 | 0.2721 | 0.1619 |

Notes: Sample size is 50,060 choice occasions involving 10,012 different households. 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.

As discussed in Section 2, the behavior of the marginal impact of a change in net expenditure, $y - p_j$, on utility is a crucial determinant of both welfare effects and pass-through of tax and cost shocks. In Figure 2 we show how the mean conditional marginal utility of income varies with $y - p_j$. Panel A shows estimates for the polynomial utility, spline utility and linear utility specifications, panel B shows numbers for the preference shifter specification and panel C focuses on the log utility specification.

The estimates of the spline utility specification show that the conditional marginal utility of income is a decreasing function of net expenditure and that over most of the domain (excluding the very bottom and top of the expenditure distribution), the function is convex. The cubic polynomial utility specification captures this shape well. 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 of income to be an increasing function of $y - p_j$. Figure 2 makes clear that in our application we do not find evidence of log-convex household demands. The linear utility model constrains the marginal utility of income to be constant and uncorrelated with consumer expenditure. This restriction is clearly not supported by the data.

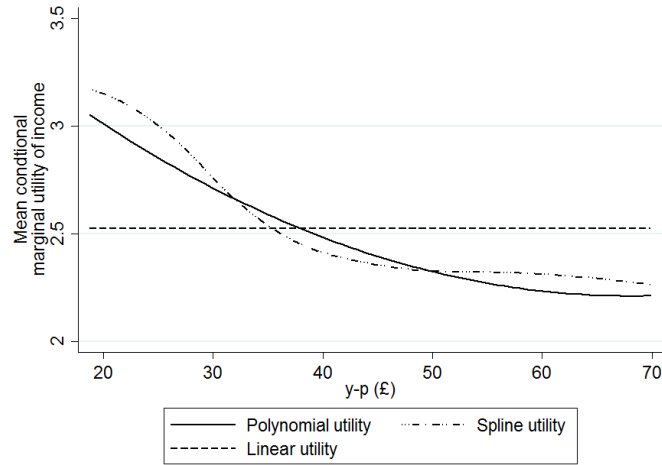
Like the linear utility specification, the preference shifter specification imposes that the conditional marginal utility of income is constant for a given household. However, it does allow the parameter to shift linearly across households based on their total expenditure, y . Panel B of Figure 2 shows that the specification does, to some extent, capture the fact that households with higher total expenditure have a lower conditional marginal utility of income. 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 polynomial specification.

The log utility specification, shown on panel C of Figure 2, yields an estimate of the conditional marginal utility of income that decreases convexly, but the function is shifted vertically downwards compared to the function implied by the polynomial specification (also shown on the graph). In principle this could reflect mis-specification of both the spline and polynomial utility models, 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 of income function. To test whether this is indeed the case we re-estimated the model specifying utility as a *third order polynomial in the log of $y - p_j$* (denoted polynomial-log utility in the figure). This model, which is more general and nests the log utility specification, yields an estimate of the conditional marginal utility of income that is very similar to the polynomial utility specification.

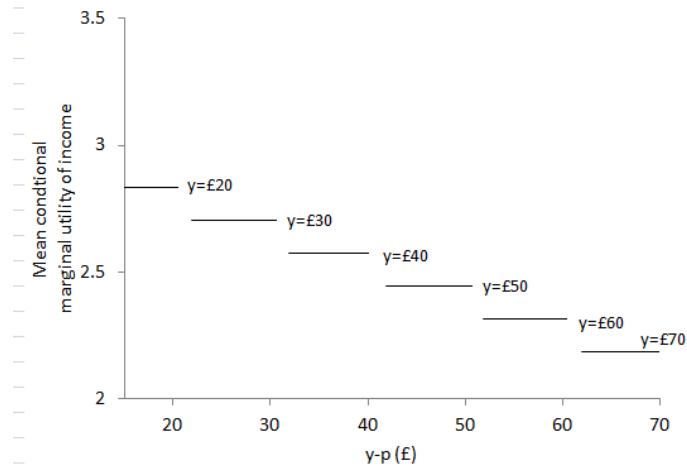
Our counterfactual analysis requires solution of a series of nonlinear first order conditions. Using estimates from the spline utility model results in relatively slow, and in some cases unstable, computations. As our baseline model we therefore proceed with the polynomial specification. It is clear from panel C of Figure 2 that, in our application, the log utility specification does a very poor job of replicating the shape of the conditional marginal utility of income 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 2: *Conditional marginal utility of income*

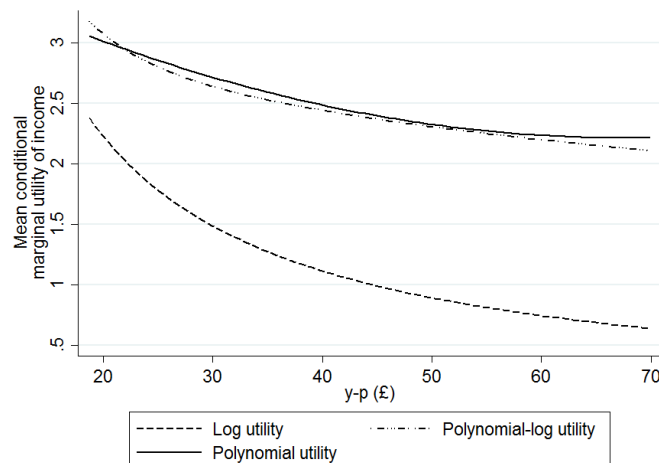
A) Spline, polynomial and linear utility



B) Preference shifter



C) Log utility



Notes: Lines shows mean conditional marginal utility of income after integrating out the random coefficients.

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 polynomial utility, linear utility and preference shifter models all yield market level price elasticity and marginal cost estimates that are very similar.⁷ 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 level elasticities are sensitive to whether we model income effects 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 we report the mean household level own price elasticity under each specification, and we report the average deviation from the mean own price elasticity for households in each quartile of the total expenditure distribution. The table also contains 95% confidence intervals.⁸ In Figure 3 we plot how household level own price elasticities vary with total 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 expenditure distribution. The polynomial utility specification results indicate that households with low expenditure are the most price sensitive; households in the bottom quartile of the expenditure distribution, on average, have an own price elasticity 0.27 below the mean and households in the top quartile, on average, have an own price elasticity 0.21 above the mean. The linear utility

⁷See Web Appendix.

⁸We 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).

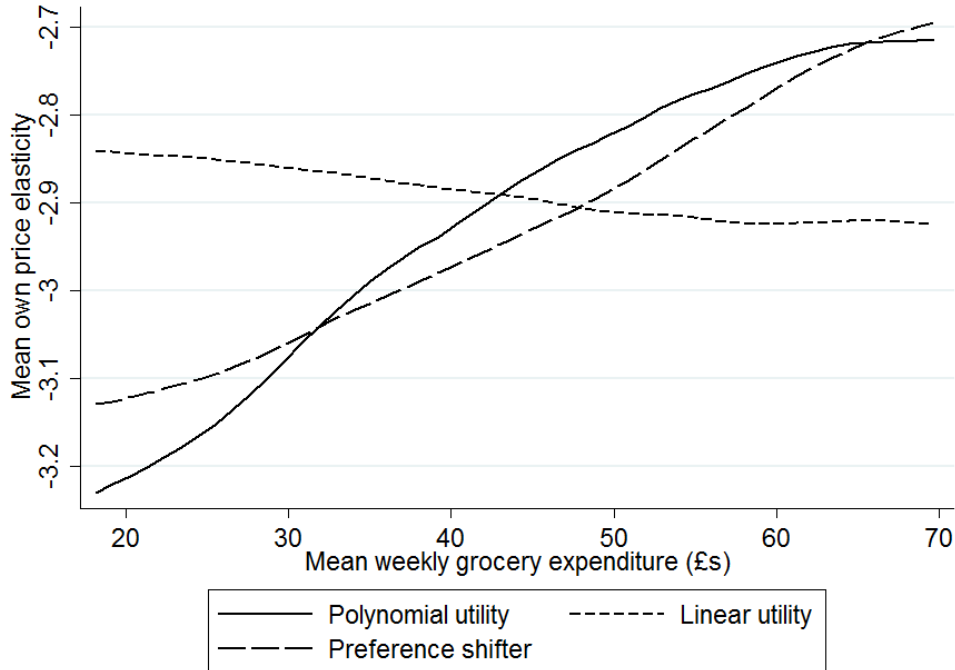
model completely fails to capture the variation in price sensitivity across the expenditure distribution, which is not surprising since expenditure plays no role in determining patterns of demand in this specification. The preference shifter specification does predict falling price sensitivity across the expenditure distribution, but it fails to capture the concavity in the relationship, underestimating price sensitivity at the bottom of the expenditure distribution and overestimating it in the center.

Table 4: *Household own price elasticity*

| Specification | Mean own price elasticity | Average deviation from mean own price elasticity for quartile of expenditure distribution: | | | |
|--------------------|---------------------------|--|-------------------------|-------------------------|-------------------------|
| | | 1 | 2 | 3 | 4 |
| Polynomial utility | -2.94 [-3.04, -2.79] | -0.27 [-0.30, -0.25] | -0.04 [-0.05, -0.03] | 0.10 [0.09, 0.11] | 0.21 [0.18, 0.23] |
| Linear utility | -2.89 [-2.97, -2.75] | 0.04 [0.03, 0.04] | 0.01 [0.02, 0.02] | -0.02 [-0.02, -0.01] | -0.04 [-0.04, -0.03] |
| Preference shifter | -2.93 [-3.06, -2.79] | -0.19 [-0.22, -0.18] | -0.07 [-0.08, -0.06] | 0.03 [0.03, 0.04] | 0.23 [0.21, 0.27] |

Notes: 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. 95% confidence intervals are shown in brackets.

Figure 3: *Variation in own price elasticities with expenditure*



Notes: 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.

4.7 Counterfactual results

To illustrate how assumptions about the marginal utility of income may affect conclusions about the impact on market equilibria and the welfare effects of policy reform we simulate the effect of an excise tax that is proportional to the saturated fat content of the product (see Section 4.4). We select the level of the tax that generates a 25% fall in purchases of saturated fat in the case of no firm pricing response (i.e. in the case of 100% pass-through).

Figure 4 is a scatter plot, at the product level, that shows how tax pass-through is related to a product's total saturated fat content. We plot the numbers for the polynomial utility specification and for three alternative specifications - the linear utility and preference shifter specifications and a simple logit specification with linear utility and with no consumer level heterogeneity.

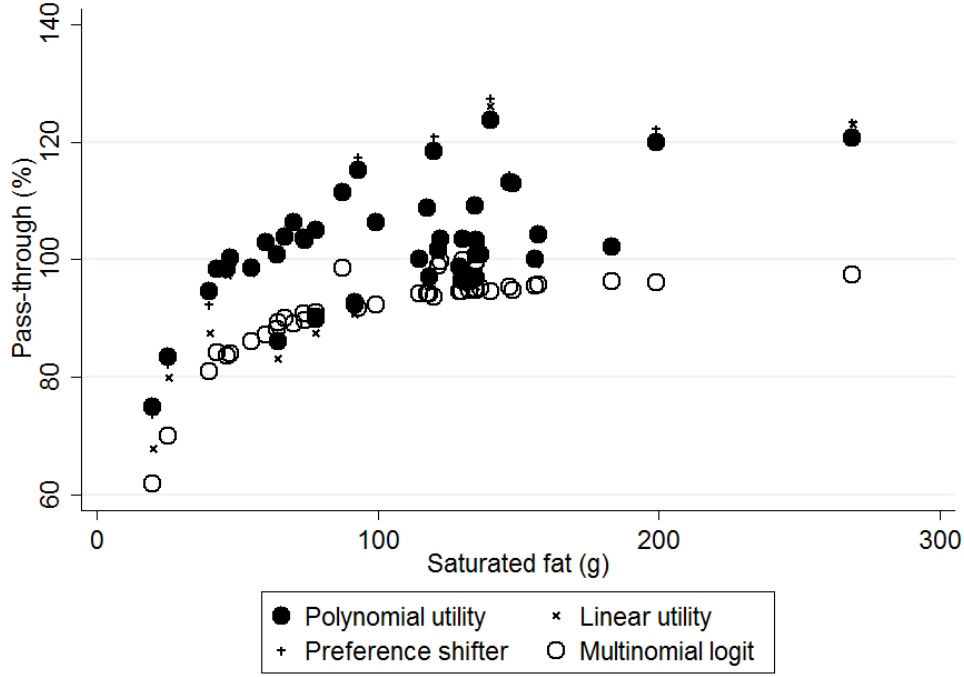
For the polynomial utility specification, across all products in the market, average pass-through of the tax to consumer prices is 103%. Therefore, on average the model predicts that prices will move close to one-to-one with the excise tax. This average masks a considerable degree of heterogeneity across products. Figure 4 shows that products with higher saturated fat contents 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 $y - p_j$ over commonly used specifications is that it relaxes restrictions on the curvature of log market demand through allowing for more flexibility in the curvature of log household demands. This flexibility allows one to test empirically whether market demand is log-concave or not. Figure 4 shows that, in our application, the alternative more restrictive linear utility specification actually yields pass-through results that are very similar to those found by the polynomial utility specification. This is also true for the preference shifter model. In this market, this suggests that the curvature restrictions placed

on household level demands (e.g. log-concavity) when utility is linear in $y - p_j$ are not overly restrictive. Together these results provide empirical evidence that market demand is log-concave in this market. If we had only estimated the linear utility model, 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 polynomial utility, linear utility and preference shifter specifications we allow for the possibility that the variance is non-zero through the inclusion of unobserved preference heterogeneity (through random coefficients). In addition, the preference shifter and polynomial models also allow for positive variance through the inclusion of 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 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 modeling pass-through.

Figure 4: *Tax pass-through across products*



Notes: For each product in each market with 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 5 we report average compensating variation estimated using each model specification. These numbers can be interpreted as the monetary payment (per year) the average household would require to be indifferent to the change in tax policy. All three models predict average compensating variation of around £2.

Columns two to five of Table 5 show the average deviation from mean compensating variation for households in each quartile of the expenditure distribution. Figure 5 shows graphically how compensating variation varies with total expenditure. All model specifications suggest compensating variation is increasing in mean weekly grocery expenditure. For the linear model the increase is comparatively small and is driven by compensating variation being related to household characteristics that are correlated with total expenditure (as the latter drops out of the model). The polynomial utility model suggests that the relationship between compensating variation and total expenditure is much stronger; on average house-

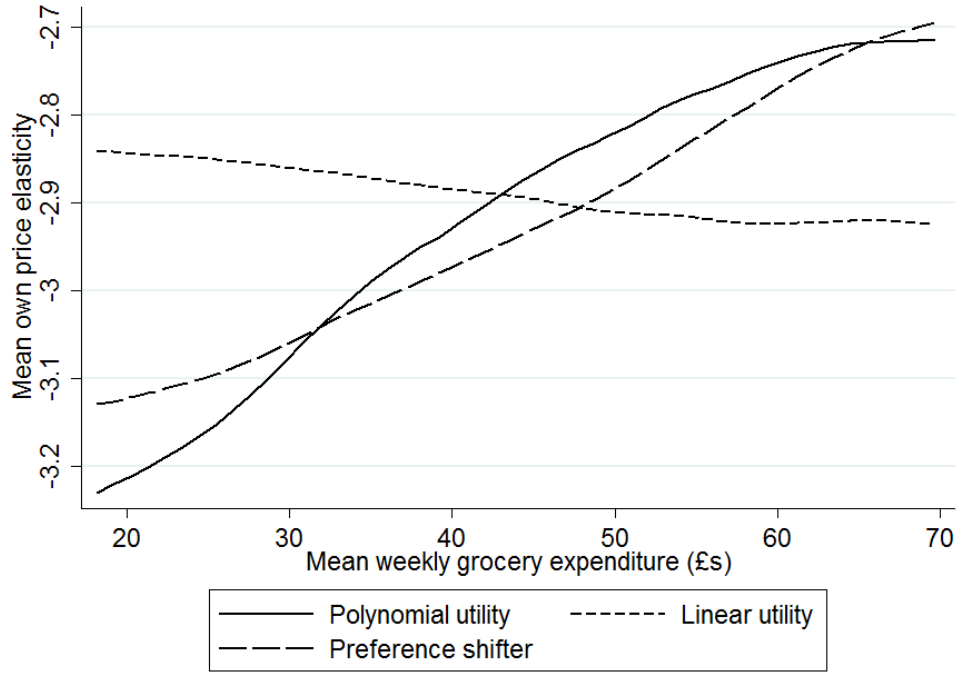
holds in the bottom quartile of the expenditure distribution have compensating variation of £0.74 below average and household in the top quartile, on average, have compensating variation £0.67 above average. Households towards the bottom of the expenditure distribution both purchase less butter and margarine and are more willing to switch between alternatives in response to a price change, leading them to be less badly affected in absolute terms than households with higher expenditure. The preference shifter model also predicts a positive relationship between a household's expenditure and compensating variation, however, it fails to capture the concavity of the relationship and overestimates compensating variation at the bottom of the expenditure distribution and underestimates it towards the center.

Table 5: *Compensating variation from tax*

| Specification | Mean compensating variation | Average deviation from mean compensating for quartile of expenditure distribution: | | | |
|--------------------|-----------------------------|--|-------------------------|----------------------|----------------------|
| | | 1 | 2 | 3 | 4 |
| Polynomial utility | 2.03 [1.95, 2.18] | -0.74 [-0.81, -0.69] | -0.19 [-0.23, -0.16] | 0.26 [0.23, 0.30] | 0.67 [0.60, 0.75] |
| Linear utility | 2.08 [1.98, 2.22] | -0.24 [-0.26, -0.21] | -0.05 [-0.05, -0.04] | 0.04 [0.04, 0.05] | 0.23 [0.21, 0.26] |
| Preference shifter | 2.05 [1.91, 2.20] | -0.63 [-0.70, -0.56] | -0.28 [-0.30, -0.23] | 0.09 [0.10, 0.12] | 0.76 [0.70, 0.87] |

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

Figure 5: *Variation in compensating variation from tax with expenditure*



Notes: Figure shows local polynomial regression of how compensating variation from tax varies with households' mean weekly grocery expenditure.

The distributional results from the preference shifter model differ from those from the polynomial utility specification because the preference shifter model does not allow enough

flexibility in the way in which y enters to fully recover how purchase patterns vary with total expenditure.

It is straightforward to demonstrate this empirically, and at the same time suggest a modification to the preference shifter model that allows it to recover the full distributional consequence 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 income 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 linearized utility model associated with equation (4.7) yields results, including distributional effects, which are very close to those from the cubic polynomial 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 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 6 and 7 illustrate for both the household level elasticities and compensating variation.

Table 6: *Mean own price elasticity: Polynomial and linearized utility*

| Specification | Mean own price elasticity | Average deviation from mean own price elasticity for quartile of expenditure distribution: | | | |
|--------------------|---------------------------|--|----------------|--------------|--------------|
| | | 1 | 2 | 3 | 4 |
| Polynomial utility | -2.94 | -0.27 | -0.04 | 0.10 | 0.21 |
| | [-3.04, -2.79] | [-0.30, -0.25] | [-0.05, -0.03] | [0.09, 0.11] | [0.18, 0.23] |
| Linearized utility | -2.94 | -0.26 | -0.03 | 0.11 | 0.20 |
| | [-3.04, -2.79] | [-0.30, -0.25] | [-0.05, -0.02] | [0.09, 0.12] | [0.17, 0.23] |

Notes: 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. 95% confidence intervals are shown in brackets.

Table 7: *Compensating variation from tax: Polynomial and linearized utility*

| Specification | Mean compensating variation | Average deviation from mean compensating variation for quartile of expenditure distribution: | | | |
|--------------------|-----------------------------|--|----------------|--------------|--------------|
| | | 1 | 2 | 3 | 4 |
| Polynomial utility | 2.03 | -0.74 | -0.19 | 0.26 | 0.67 |
| | [1.95, 2.18] | [-0.81, -0.69] | [-0.23, -0.16] | [0.23, 0.30] | [0.60, 0.75] |
| Linearized utility | 2.04 | -0.76 | -0.21 | 0.25 | 0.66 |
| | [1.95, 2.19] | [-0.82, -0.70] | [-0.23, -0.16] | [0.23, 0.31] | [0.60, 0.75] |

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

In an appendix we shows that if we alternatively consider an ad valorem tax, of the form described in Section 4.4, we find that this tax is under-shifted. Our conclusions regarding the effects of not modelling income effects flexibly and in a theoretically rigorous way remain very similar.

5 Conclusion

In this paper we have explored the importance of relaxing restrictions commonly placed on the marginal utility of income in logit demand models. By far the two most common

approaches are either to assume that the marginal utility of income is constant for a given consumer, but to allow it to vary cross-sectionally with demographics including consumer income, or to model income effects by assuming utility is linear in the log of all spending outside the market currently under focus. Both of these approaches heavily constrain income effects (ruling them out in the first case) and unduly restrict demand curvature. Imposing these restrictions prevent the data from providing evidence as to the true shape of the demand curve.

Specifying consumer level utility in the form $U_j = f(y - p_j) + g(\mathbf{x}_j) + \epsilon_j$ for some flexible nonlinear (e.g. polynomial) function $f(\cdot)$ offers three advantages. Firstly, it allows the model to capture any income effects induced by the policy counterfactual under consideration. Secondly, it allows for more flexibility in the curvature of consumer level demands. Thirdly it allows for a richer relationship between expenditure, demand, and welfare.

To explore the empirical importance of these restrictions we consider an application to the UK butter and margarine market. This product category comprises a small fraction of households' budgets and is a category for which flexible modeling of the marginal utility of income may *a priori* not seem to be of first order importance. Yet we show that results from a flexible model differ from results from standard models in important ways. In the case of the log utility specification, it is clear that the shape imposed on the conditional marginal utility of income is too restrictive leading the log utility model to yield implausible predictions.

The commonly used but ad hoc preference shifter model does a good job of replicating market level average elasticities, marginal costs, pass-through and consumer welfare but is less successful in recovering distributional aspects of demand and welfare effects. If researchers are interested in the distributional consequences of reforms that result in price changes that are small relative to total income, they should consider either flexibly incorporating income effects, or using more flexible models of preference shifting.

In our application the marginal utility of income 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 results of the polynomial utility specification.

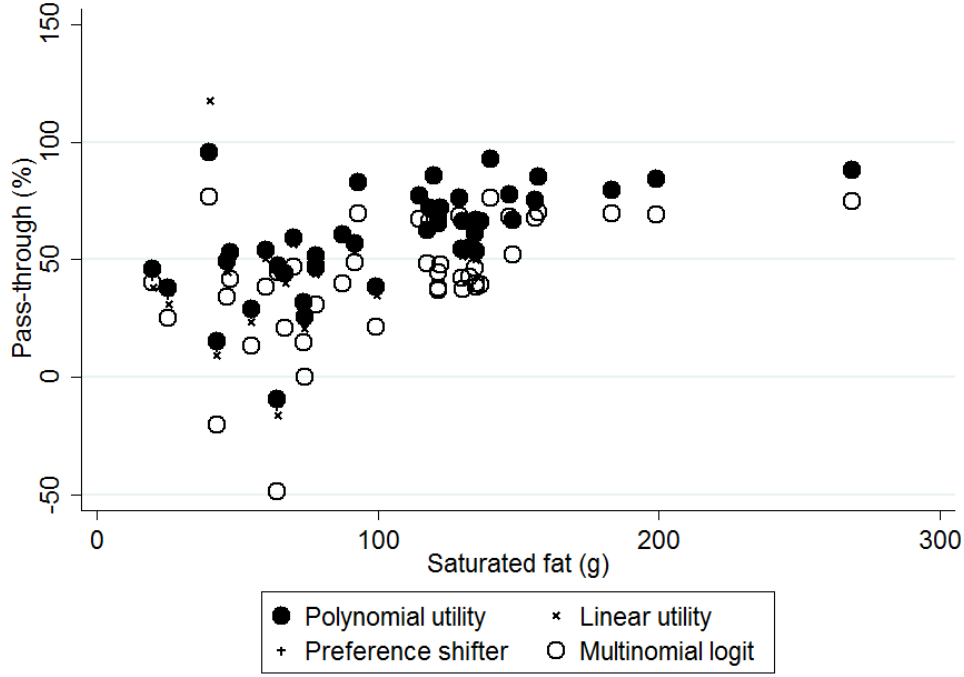
In applications to product categories comprising large shares of consumers' budgets, flexibly modeling income effects is likely to be even more important than in our application. In such markets, price changes are more likely to be large enough to induce significant income effects. In applications involving 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 of income and should be tested against more flexible alternative specifications.

A Appendix: An ad valorem tax

In Section 4.7 we report results from simulating an excise tax that is proportional to products' saturated fat contents and that generates a 25% fall in purchases of saturated fat in the case of no firm pricing response (i.e. in the case of 100% pass-through). Here we report results from simulating an ad valorem tax that is proportional to products' saturated fat content and that generates a 25% fall in purchases of saturated fat under 100% pass-through. Figure 6 shows the patterns of pass-through across products for each model specification (the ad valorem tax analog of Figure 4). In contrast to the excise tax, the ad valorem tax is under-shifted to final prices - average pass-through under the polynomial utility model is 58%. However, as with the excise tax, the polynomial utility, linear utility and preference shifter models generate the same pattern of pass-through across products and the multinomial logit model generates pass-through that, on average, is lower (43% on average).

Table 8 describes compensating variation from the ad valorem tax (the ad valorem tax analog of Table 5). As the ad valorem tax is under-shifted to consumer prices, compensating variation from the tax is less than for the excise tax. In common with the excise tax, the preference shifter and linear utility models fail to fully replicate how compensating variation varies across the expenditure distribution under the polynomial utility specification.

Figure 6: *Ad valorem tax pass-through across products*



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

Table 8: *Compensating variation from ad valorem tax*

| Specification | Mean compensating variation | Average deviation from mean compensating for quartile of expenditure distribution: | | | |
|--------------------|-----------------------------|--|-------------------------|----------------------|----------------------|
| | | 1 | 2 | 3 | 4 |
| Polynomial utility | 1.33 [1.29, 1.39] | -0.59 [-0.64, -0.56] | -0.18 [-0.21, -0.15] | 0.21 [0.18, 0.24] | 0.57 [0.51, 0.64] |
| Linear utility | 1.35 [1.29, 1.43] | -0.19 [-0.19, -0.16] | -0.04 [-0.04, -0.03] | 0.03 [0.03, 0.04] | 0.18 [0.16, 0.20] |
| Preference shifter | 1.34 [1.27, 1.41] | -0.52 [-0.56, -0.46] | -0.25 [-0.26, -0.22] | 0.06 [0.07, 0.09] | 0.65 [0.61, 0.74] |

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

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.
- 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.
- Fabinger, M. and E. G. Weyl (2014). A tractable approach to pass-through patterns with applications to international trade. *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.
- 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.

- Train, K. E. (2003). *Discrete Choice Methods with Simulation*. Cambridge University Press.
- 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.

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

Rachel Griffith, Lars Nesheim and Martin O'Connell

June 4, 2015

In this appendix we include estimated market price elasticities and marginal costs for the polynomial utility, linear utility and preference shifter specifications. Table 1 presents a matrix of average market own and cross price elasticities for the 10 products with the highest market share. It contains the matrix for each of the three model specifications. The numbers show: 1) Demand for all products is elastic, with own price elasticities ranging from -1.7 to -5.0. 2) Cross-price elasticities exhibit a high degree of variation, showing estimates are far from those from a conditional logit (in which there would be no within column variation in cross price elasticities). The cross-price elasticities also indicate a much higher degree of substitution within the butter products (Ar: Anchor NZ 500g, Ar: Lurpak spread ss 500g and Ar: Lurpak light ss 500g) than between them and the margarine products. 3) The three models yield similar estimates for market own and cross price elasticities. This contrasts with their predictions for household level elasticities which differ (see Table 4 and Figure 3 of the main paper).

Table 2 presents the mean marginal cost estimates for the 10 largest market share products. They are based on the assumption that firms compete in a Nash-Bertrand game and therefore are a functions of the market level price elasticities and the ownership structure of products. Given the similarities in market elasticities between the three model specifications, it is not surprising that the models generate a similar set of marginal costs. Margins are estimated to be lower for the butter products than for the margarine products.

Table 1: Own and cross price elasticities

| Polynomial utility | | | | | | | | | | | | | |
|---------------------------|------------------------------|-----------------------------|---------------------------|------------------------------|---------------------------|-------------------------|-------------------|--------------------|--------------------------|--------------------|-------------------|-------------------|-------------------|
| Ar: Anchor NZ 500g | Ar: Lurpak spread ss 500g | Ar: Lurpak light ss 500g | DC: Clover spread 500g | DC: Utterly Butterly 500g | Un: Flora buttery 500g | Un: Flora light 500g | Un: Flora 500g | Un: ICBINB 500g | Un: ICBINB Light 500g | Un: ICBINB 500g | Un: Stork 500g | Outside option | Outside option |
| Ar: Anchor NZ 500g | -4.7723 | 0.1465 | 0.0333 | 0.0357 | 0.0344 | 0.0340 | 0.0358 | 0.0356 | 0.0356 | 0.0361 | 0.0361 | 0.0364 | 0.0364 |
| Ar: Lurpak spread ss 500g | 0.2470 | -4.9456 | 0.0570 | 0.0593 | 0.0600 | 0.0585 | 0.0605 | 0.0602 | 0.0602 | 0.0603 | 0.0603 | 0.0597 | 0.0597 |
| Ar: Lurpak light ss 500g | 0.1757 | 0.1852 | 0.0412 | 0.0431 | 0.0436 | 0.0425 | 0.0439 | 0.0437 | 0.0437 | 0.0437 | 0.0437 | 0.0429 | 0.0429 |
| DC: Clover spread 500g | 0.0346 | 0.0356 | -2.4253 | 0.0516 | 0.0524 | 0.0538 | 0.0493 | 0.0495 | 0.0495 | 0.0505 | 0.0505 | 0.0423 | 0.0423 |
| DC: Utterly Butterly 500g | 0.0239 | 0.0236 | 0.0331 | -1.7729 | 0.0324 | 0.0331 | 0.0330 | 0.0328 | 0.0328 | 0.0331 | 0.0331 | 0.0304 | 0.0304 |
| Un: Flora butterly 500g | 0.0412 | 0.0429 | 0.0603 | 0.0576 | -2.1381 | 0.0604 | 0.0598 | 0.0600 | 0.0600 | 0.0584 | 0.0584 | 0.0517 | 0.0517 |
| Un: Flora light 500g | 0.0270 | 0.0279 | 0.0411 | 0.0395 | 0.0403 | -2.4811 | 0.0403 | 0.0404 | 0.0404 | 0.0390 | 0.0390 | 0.0320 | 0.0320 |
| Un: ICBINB 500g | 0.0328 | 0.0332 | 0.0443 | 0.0461 | 0.0460 | 0.0464 | -1.8513 | 0.0477 | 0.0477 | 0.0458 | 0.0458 | 0.0417 | 0.0417 |
| Un: ICBINB light 500g | 0.0166 | 0.0167 | 0.0225 | 0.0233 | 0.0234 | 0.0236 | 0.0242 | -1.9077 | -1.9077 | 0.0232 | 0.0232 | 0.0210 | 0.0210 |
| Un: Stork 500g | 0.0131 | 0.0131 | 0.0176 | 0.0180 | 0.0177 | 0.0176 | 0.0180 | 0.0180 | 0.0180 | -1.6520 | -1.6520 | 0.0173 | 0.0173 |
| Linear utility | | | | | | | | | | | | | |
| Ar: Anchor NZ 500g | Ar: Lurpak spread ss 500g | Ar: Lurpak light ss 500g | DC: Clover spread 500g | DC: Utterly Butterly 500g | Un: Flora buttery 500g | Un: Flora light 500g | Un: Flora 500g | Un: ICBINB 500g | Un: ICBINB Light 500g | Un: ICBINB 500g | Un: Stork 500g | Outside option | Outside option |
| Ar: Anchor NZ 500g | -4.7133 | 0.1452 | 0.0335 | 0.0363 | 0.0351 | 0.0342 | 0.0366 | 0.0364 | 0.0364 | 0.0369 | 0.0369 | 0.0375 | 0.0375 |
| Ar: Lurpak spread ss 500g | 0.2484 | -4.8383 | 0.0581 | 0.0611 | 0.0618 | 0.0597 | 0.0626 | 0.0623 | 0.0623 | 0.0624 | 0.0624 | 0.0622 | 0.0622 |
| Ar: Lurpak light ss 500g | 0.1772 | 0.1873 | 0.0421 | 0.0445 | 0.0450 | 0.0435 | 0.0453 | 0.0453 | 0.0453 | 0.0452 | 0.0452 | 0.0448 | 0.0448 |
| DC: Clover spread 500g | 0.0341 | 0.0350 | -2.3811 | 0.0518 | 0.0525 | 0.0536 | 0.0496 | 0.0498 | 0.0498 | 0.0508 | 0.0508 | 0.0423 | 0.0423 |
| DC: Utterly Butterly 500g | 0.0235 | 0.0232 | 0.0328 | -1.7471 | 0.0322 | 0.0328 | 0.0330 | 0.0328 | 0.0328 | 0.0332 | 0.0332 | 0.0300 | 0.0300 |
| Un: Flora butterly 500g | 0.0407 | 0.0423 | 0.0599 | 0.0577 | -2.0934 | 0.0601 | 0.0598 | 0.0601 | 0.0601 | 0.0586 | 0.0586 | 0.0511 | 0.0511 |
| Un: Flora light 500g | 0.0267 | 0.0275 | 0.0410 | 0.0397 | 0.0404 | -2.4320 | 0.0405 | 0.0406 | 0.0406 | 0.0392 | 0.0392 | 0.0320 | 0.0320 |
| Un: ICBINB 500g | 0.0324 | 0.0326 | 0.0440 | 0.0461 | 0.0457 | 0.0460 | -1.8185 | 0.0477 | 0.0477 | 0.0459 | 0.0459 | 0.0410 | 0.0410 |
| Un: ICBINB light 500g | 0.0163 | 0.0165 | 0.0224 | 0.0233 | 0.0232 | 0.0234 | 0.0241 | -1.8736 | -1.8736 | 0.0232 | 0.0232 | 0.0206 | 0.0206 |
| Un: Stork 500g | 0.0128 | 0.0128 | 0.0173 | 0.0179 | 0.0175 | 0.0174 | 0.0178 | 0.0178 | 0.0178 | -1.6266 | -1.6266 | 0.0169 | 0.0169 |
| Preference shifter | | | | | | | | | | | | | |
| Ar: Anchor NZ 500g | Ar: Lurpak spread ss 500g | Ar: Lurpak light ss 500g | DC: Clover spread 500g | DC: Utterly Butterly 500g | Un: Flora buttery 500g | Un: Flora light 500g | Un: Flora 500g | Un: ICBINB 500g | Un: ICBINB Light 500g | Un: ICBINB 500g | Un: Stork 500g | Outside option | Outside option |
| Ar: Anchor NZ 500g | -4.7609 | 0.1458 | 0.0332 | 0.0357 | 0.0344 | 0.0339 | 0.0358 | 0.0356 | 0.0356 | 0.0361 | 0.0361 | 0.0367 | 0.0367 |
| Ar: Lurpak spread ss 500g | 0.2466 | -4.9109 | 0.0570 | 0.0594 | 0.0600 | 0.0586 | 0.0606 | 0.0604 | 0.0604 | 0.0604 | 0.0604 | 0.0602 | 0.0602 |
| Ar: Lurpak light ss 500g | 0.1756 | 0.1858 | 0.0413 | 0.0432 | 0.0437 | 0.0426 | 0.0440 | 0.0438 | 0.0438 | 0.0438 | 0.0438 | 0.0433 | 0.0433 |
| DC: Clover spread 500g | 0.0347 | 0.0356 | -2.4295 | 0.0519 | 0.0526 | 0.0540 | 0.0495 | 0.0497 | 0.0497 | 0.0508 | 0.0508 | 0.0426 | 0.0426 |
| DC: Utterly Butterly 500g | 0.0239 | 0.0237 | 0.0333 | -1.7807 | 0.0326 | 0.0333 | 0.0332 | 0.0331 | 0.0331 | 0.0334 | 0.0334 | 0.0307 | 0.0307 |
| Un: Flora butterly 500g | 0.0413 | 0.0430 | 0.0606 | 0.0580 | -2.1433 | 0.0608 | 0.0602 | 0.0604 | 0.0604 | 0.0588 | 0.0588 | 0.0522 | 0.0522 |
| Un: Flora light 500g | 0.0271 | 0.0279 | 0.0412 | 0.0397 | 0.0405 | -2.4840 | 0.0405 | 0.0406 | 0.0406 | 0.0392 | 0.0392 | 0.0323 | 0.0323 |
| Un: ICBINB 500g | 0.0330 | 0.0333 | 0.0446 | 0.0464 | 0.0463 | 0.0467 | -1.8579 | 0.0481 | 0.0481 | 0.0462 | 0.0462 | 0.0421 | 0.0421 |
| Un: ICBINB light 500g | 0.0166 | 0.0168 | 0.0227 | 0.0234 | 0.0235 | 0.0237 | 0.0244 | -1.9143 | -1.9143 | 0.0234 | 0.0234 | 0.0211 | 0.0211 |
| Un: Stork 500g | 0.0131 | 0.0132 | 0.0177 | 0.0182 | 0.0179 | 0.0178 | 0.0182 | 0.0182 | 0.0182 | -1.6600 | -1.6600 | 0.0174 | 0.0174 |

Notes: Each cell contains the price elasticity of demand for the product indicated in row 1 with respect to the price of the product in column 1. Numbers are means across markets.

Table 2: *Marginal costs: top 10 market share products*

| | Price | Polynomial utility | | Linear utility | | Preference shifter | |
|---------------------------|-------|--------------------|--------|----------------|--------|--------------------|--------|
| | | Cost | Margin | Cost | Margin | Cost | Margin |
| Ar: Anchor NZ 500g | 1.99 | 1.49 | 0.25 | 1.49 | 0.26 | 1.49 | 0.25 |
| Ar: Lurpak spread ss 500g | 2.15 | 1.64 | 0.24 | 1.63 | 0.24 | 1.64 | 0.24 |
| Ar: Lurpak light ss 500g | 2.17 | 1.66 | 0.24 | 1.64 | 0.24 | 1.65 | 0.24 |
| DC: Clover spread 500g | 1.20 | 0.67 | 0.44 | 0.66 | 0.45 | 0.67 | 0.44 |
| DC: Utterly Butterly 500g | 0.80 | 0.32 | 0.60 | 0.31 | 0.61 | 0.32 | 0.60 |
| Un: Flora buttery 500g | 1.02 | 0.46 | 0.55 | 0.44 | 0.57 | 0.46 | 0.55 |
| Un: Flora light 500g | 1.22 | 0.63 | 0.48 | 0.62 | 0.50 | 0.63 | 0.48 |
| Un: ICBINB 500g | 0.85 | 0.31 | 0.64 | 0.30 | 0.65 | 0.31 | 0.63 |
| Un: ICBINB light 500g | 0.87 | 0.33 | 0.62 | 0.31 | 0.64 | 0.33 | 0.62 |
| Un: Stork 500g | 0.72 | 0.20 | 0.72 | 0.19 | 0.74 | 0.20 | 0.72 |

Notes: Margins are defined as $(p - mc)/p$. Numbers are market share weighted means.