Are Heavier Drinkers Less Responsive to Price? Quantile Regression Estimates

Bruce Hollingsworth[a], Robert Pryce[ab*], and Ian Walker[b]

[a]Division of Health Research, Lancaster University
[b]Economics Department, Lancaster University

November 2013

This is a working paper.
Please do not quote or reference this paper without the authors' permission.

Abstract

Many people in the United Kingdom drink more than the recommended level of alcohol, some drinking substantially more. This leads to large health and social costs, and price is often proposed as a tool for reducing consumption. This paper uses quantile regression methods to estimate the differential price (and income) elasticities across the drinking distribution, without resorting to endogenously splitting the sample into subgroups. The paper uses data from the UK Expenditure and Food Survey (EFS) from 2001 to 2010, and finds that heavy drinkers are less responsive to price than light or moderate drinkers, especially for drinking away from licensed premises ('off premise'). The implication is that price-based policies may have little effect in reducing consumption amongst the heaviest drinkers, who are often the cause of most health and social harms. Very high taxes that reduce heavy drinking to acceptable levels would generate very large deadweight losses on the majority of drinkers who do so in moderation. Other policy options may be more productive in reducing consumption in this group.

Keywords: Alcohol, Quantile Regression

JEL Classification: D12, I18

HESG Reference A018

*Corresponding Author. Email r.pryce@lancaster.ac.uk. Robert Pryce is a PhD student funded by an ESRC NWDTC CASE scholarship. We are grateful to the Economic and Social Data Service (now UK Data Service) for providing the data.

1 Introduction

£42.1 billion was spent on alcohol in England and Wales in 2010 (Clancy, 2011). Over 5 million hectolitres of pure alcohol was cleared by Her Majesty's Revenues and Customs, which generated over £9 billion in tax revenue (HMRC, 2012). However, the UK Government's Alcohol Strategy estimates that alcohol-related harms cost society £21 billion per annum (Home Office, 2012). Alcohol is the second largest risk factor for disease and disability in Europe (WHO, 2011) causing 2.5 million deaths per year globally, and causing 9% of deaths in those aged 15-29. Internationally alcohol is a bigger killer than tobacco or high blood pressure (WHO, 2011).

The World Health Organisation advocates tax increases as one means of reducing consumption (WHO, 2011). Not surprisingly, the price elasticity of demand for alcohol has been the subject of considerable research (see Fogarty (2010) for a detailed summary). This parameter is clearly important to governments who want to model the impact of tax increases on tax receipts, as well as public health authorities who wish to reduce alcohol consumption for health and social reasons. Alcohol imposes quantitatively significant and probably non-linear external costs (see, for example, Dee (1999) or Leon and McCambridge (2006)) which a Pigouvian tax would seek to internalise (see, for example, Kenkel (1996)). However, it would be difficult to make this tax correspondingly non-linear because of an inability to prevent customers arbitraging price differentials. Minimum pricing is under consideration in several jurisdictions, which may provide a partial solution to resolving this problem. Flat taxes would likely lead to large welfare losses, as moderate drinkers would be taxed for externalities which they do not create. Indeed, it may also be the case that heavier drinkers respond differently to price than do moderate drinkers. If heavier drinkers are less responsive to a change in price, which seems likely because of the habituating nature of the product, then any tax aimed at reducing average consumption will cause even greater welfare losses than would otherwise be the case. Knowing the price responsiveness of the heaviest drinkers is therefore very important to policymakers, since it is this group who potentially cause the greatest harms.

Several papers have explored how the price elasticity varies across the drinking distribution by estimating over light, moderate and heavy drinkers separately. However, estimation using endogenously selected sub-samples according to the extent of consumption is likely to result in estimated elasticities that are biased upwards[1]. The aim of this paper is to estimate the price (and income) elasticity for alcohol across the drinking distribution, using quantile regression. This technique is relatively under-used in this context, and yet it avoids the inevitable bias when estimating using endogenously selected subgroups. The contributions of the paper are that: it provides UK evidence on the price elasticity of demand for alcohol; it uses data pooled over more than a decade which provides considerable relative price variation; and it shows statistically significant variation in the demand elasticity across the distribution of drinking, which has important implications for policy.

The results show significantly lower price elasticities for heavy drinkers. We go on to explore the consequences of raising taxes in the light of these estimates. Our conclusion is that the level of tax that is sufficient to reduce heavy drinking to acceptable levels would impose very large deadweight losses on the moderate majority.

[1] It can be shown that the price elasticity of demand will be biased upwards (in absolute terms) if selection is done by the quantity consumed. See Heckman (1979) for the seminal analysis of selection bias.

2 Estimation Issues

The demand for alcohol has been widely researched. A meta-analysis by Gallet (2007), which reviewed estimates from 132 studies, found a median own-price elasticity of -0.497. A similar meta-analysis by Wagenaar et al. (2009), of 112 studies, found a mean own-price elasticity of -0.51. However, neither of these meta-analyses is specific to the United Kingdom, and it is likely that different countries with different attitudes to alcohol will have different responses to price. Recent work by Collis et al. (2010) estimates price elasticities for the United Kingdom. It splits alcohol in to 10 different products: 5 drinks (beer, cider, wine, spirits, ready-to-drink), either on licensed premises (pubs, clubs and restaurants) or off-premise (off-licenses - supermarket and other licensed shops). They find own-price elasticity estimates ranging from -0.46 (for on-premise wine) to -1.34 (for off-premise cider). Since the cross-price elasticities between alcoholic products are often significantly negative, it is to be expected that the demand for alcohol as a whole should be less elastic to the extent that there is substitution between types within this product group.

There are many important difficulties in estimating price elasticities for alcoholic products. Many individuals do not drink, and would not do so at any likely price (abstainers); many drinkers drink with such infrequency that expenditure surveys do not capture their consumption very well (infrequency of purchase); some individuals do not drink at current prices but may at lower prices (corner solutions); and some individuals have expenditures that do not match their consumption, even in the long run, because buying and consuming may be done by different people, especially within households (intra-household redistribution). A Tobit specification is a restrictive solution: it implies that participation elasticities are proportional to elasticities for positive expenditure that makes it an appropriate solution only in the classical corner solutions case but not otherwise. Dealing with abstention could be achieved with a censored regression framework but requires an exclusion restriction for identification purposes - there has to be some variable that contributes to the determination of participation that does not affect expenditure conditional on participation. It is difficult to imagine what such a variable might be. Combining abstention with corner solutions could be achieved with a double hurdle model (see, for example, Jones' (1989) paper on tobacco demand) but an exclusion restriction is still required. Keen (1986) has provided a solution to the problem of zeroes induced by infrequency of purchase, that replaces observed expenditure levels with expectations based on the proportion of zeroes in the data. Combining infrequency with abstention/corner solutions has not, to our knowledge, been done, and there have been no studies that attempt to resolve the mismatch between consumption and expenditure, apart from infrequency.

Another important problem is that the price at which alcohol is purchased is likely to be endogenous: in particular heavy drinkers have the greatest incentive to search for the lowest price and to buy in bulk at a discount (price endogeneity). This will bias elasticity estimates towards zero. Moreover, price is determined, in part, by quality; and quality and quantity may well be correlated because heavy drinkers are likely to be drinking lower price alcohol. Indeed, it is usually not possible to observe the price at which alcohol is purchased and this has to be estimated in some way: leading to the potential for endogeneity in the imputed price. For example, several studies have divided expenditure by quantity purchased leading to a form of measurement error (known as division bias, because the price variable is obtained by dividing by the dependent variable)[2]. A further problem arises if quantity and/or expenditure is under-recorded and this under-recording is correlated with true consumption. If the correlation between reported and true were zero then this systematic mismeasurement would simply affect the constant term in a double log specification - although even if the proportion of under-recording varied randomly

[2] In all cases that we have noted in the literature this potentially serious source of endogeneity is ignored. This is surprising since it cane be avoided by a simple reparameterisation, at least in the double log specification that is the dominant specification in the literature.

across individuals this would bias the price elasticity (towards zero) if price were computed by dividing reported expenditure by reported quantity except in the unlikely event that the under-recording in each were the same. However, if the extent of under-recording is increasing with true consumption, as seems likely, then the price elasticity will be biased downwards even if there is no division bias arising from how price is computed.

The recent research that has used prices derived from the micro data by dividing expenditure by quantity is vulnerable to bias associated with unobserved heterogeneity in quality to the extent that quality affects price, as well as division bias. The direction of bias is complicated and will depend on the price elasticity of quantity, and of quality, the effect of quality on price, and the extent to which quality and quantity are normal goods. Resolving this issue is likely to be complicated and the alternative of using a price index derived from grouped or aggregate data has much to recommend it.

Aggregation across alcoholic products into total units of alcohol may be possible, but only under restrictive conditions on preferences and/or prices (see, for example, Muellbauer (1976)). Existing research has never, to our knowledge, considered this as an issue yet if the aggregation conditions are not satisfied the resulting estimates have no status as elasticities. Nonetheless, aggregation across products may be desirable to reduce the problem arising from having many zero expenditures in the data. Aggregation across household members may be required if only household data is available and, in any event, may be desirable because of the likely mismatch between an individual's consumption and their expenditures that could result in many zeroes in the individual expenditure data. Household level data will need to allow for variation in household composition in the modeling. There are some further issues associated with the specification of the demand relationship. For example, it is probably inevitable that some parametric assumption will be needed to facilitate such estimation[3]. A popular choice is a double log specification because the log price coefficient can then be interpreted as an elasticity. However, this would not be an appropriate specification to apply to a dataset containing drinkers and non-drinkers since it would not enable incorporation of the zeroes, and many studies therefore use semi-log models or model alcohol expenditure or the share of alcohol in total expenditure[4]. While a small proportion of the existing work takes a systems approach to modeling the demand for several alcoholic drinks and in different locations to take account of cross-price effects, none of the existing work allows for cross-price elasticities with other, non-alcoholic, products[5]. There are further practical problems from using estimated elasticities for policy purposes. For example, the alcoholic content of drinks are usually not known so converting from observed behavior into health harms can only be done with error and this error is correlated with the likely harms.

[3]Semi-parametric methods could conceivably be used to allow for a flexible price specification with parametric intercept shifts to account for differences in characteristics. However, it seems unlikely that there would be sufficiently continuous price variation to do this successfully. Moreover, it would be difficult to apply such semi-parametric methods with endogenous prices and, more generally, in non-linear models such as the Tobit

[4]The Almost Ideal Demand System of Deaton and Muellbauer (1980), which has the share of total expenditure as the dependent variable and log price and log income (and, sometimes, its square) is widely regarded as being a flexible model that provides a good fit to household spending data. See Blundell et al. (1999) for an example in the context of fuel spending patterns. Of course, it makes little sense to use a budget share approach within the context of quantile regression, since rich heavy-drinking households may have a smaller budget share than poor light-drinking households.

[5]An exception is Symons and Walker (1990) which considers the demand for beer, wines and spirits within a broader demand system and identifies statistically significant cross-price effects. Importantly, it seems likely that tobacco and alcohol are complements, and there was a significant cross elasticity with respect to the price of domestic fuel.

3 Existing Literature

There are five studies in the large existing literature that are particularly notable and relevant to this research. Collis et al (2010) is a careful analysis of five years of the UK household expenditure surveys - the Expenditure and Food Survey (EFS). The data is well known to under-record alcohol spending, but is very detailed. They consider five types of drink, both in licensed premises and in licensed shops. They use data at the household level. They use a Tobit model to allow for zero consumption, and the dependent variable is quantity of alcohol in millilitres at the household level. Price is generated individually for each household by dividing expenditure by quantity. Zeroes are retained for fear of generating bias from dropping them. Non-purchasing households are assumed to face a price based on the year, region and household size averages of drinkers - this imputed price is applied, inconsistently, to the zeroes but not to those making positive expenditures. There are therefore (at least) two potential issues with the Collis et al (2010) study. The first is that prices are generated for each individual household by dividing expenditure by quantity, which generates an important source of endogeneity bias. If heavier drinking households purchase cheaper alcohol, then elasticity estimates will be biased upwards (more elastic). Secondly, the study assumes that all non-purchasing households are at corner solutions, meaning that they do not purchase alcohol simply because the price is too high, relative to other goods and income. If zero expenditure arises from other reasons, as well as corner solutions, then the price elasticity that comes out of the Tobit specification will be biased towards zero. The Tobit model imposes the constraint that the elasticity of participation is proportionate to the elasticity conditional on participation - for all explanatory variables. This is a highly restrictive assumption and one that is unlikely to be true. Of course, the opposite extreme would be to drop all zeroes from the data. Since they consider 10 different alcoholic products this would generate a large majority of households with at least one zero. The Tobit is too restrictive a specification - the participation decision needs to be modelled separately from the consumption decision by participants. Jones (1989) considers a double hurdle model in the context of smoking and Jones and Yen (2000) propose an extension that is particularly suited to a logarithmic specification. However, none of the existing literature has provided a plausible solution to the non-parametric identification problem. This requires a valid exclusion restriction which is particularly difficult in this context.[6]

The second notable study is Purshouse et al (2010)[7]. Like Collis et al (2010), this also uses the UK household expenditure survey. This also uses the same five years of EFS data and looks at four products divided into two price levels and in two locations; and it models expenditures at the individual level rather than the household level. The motivation is to relate individual consumption to individual harms, although there is no attempt to validate the assumption that consumption equals expenditure. The observations are divided into moderate and heavy drinkers and no attempt is made for the likely selection bias that this induces. Moreover, it is a double log specification and there is no attempt to model the zeroes in the data[8].

Saffer et al. (2012) is our third relevant study. It uses information on US youths that have been followed through time and are aged between 17 and 28. Since this is a US dataset, a large proportion of respondents will not be legally able to purchase alcohol and being below the minimum age is included as an independent variable. Saffer et al. (2012) use both a finite mixture model (FMM) and a quantile regression (QR) model. The first specifies a number of sub-groups, each with its own model of behaviour and allows the data to determine the probability that each observation belongs to each sub-group as well

[6] An alternative solution is pursued in Farrell and Walker (1999) which uses a probit to model participation and a censored least absolute deviation (CLAD) specification to model the consumption of participants in the context of a model of gambling.
[7] The econometric analysis can be found in the web appendix to their paper.
[8] In addition, no standard errors are reported.

as the parameters of each model. The second is a generalisation of the median regression (or least absolute deviation, LAD) which is analogous to conventional least squares but minimises the sum of the absolute differences between actual and predicted, rather than the squared differences. QR generalizes LAD by minimizing weighted absolute deviations. Neither FMM or quantile model requires selection by the dependent variable. QR has the considerable advantage that it uses all of the data to estimate each quantile regression and simply varies the weights.

The dependent variable used is the number of drinks per month, calculated as the product of the number of drinking days per month and the usual number of drinks consumed per day. Of course, there may be uncertainty as to what a "drink" is, which could lead to measurement error in the dependent variable which will bias (upward) the estimated standard errors but not the coefficients themselves. There is no on-premise/off-premise split. The model is double log and the zeroes are dropped with no correction. The FMM specification predicts that there are two sub-groups, moderate and heavy drinkers, and moderate drinkers have a price elasticity of -0.49, whilst the price elasticity for heavy drinkers is not significantly different to zero. Their QR results suggest that only light drinkers are significantly responsive to price at a 5% level: their price elasticities become more inelastic along the drinking distribution (except for the 60th percentile) and after the 60th percentile, their price elasticity is not statistically significantly different from zero. In contrast, Saffer et al. (2012) suggests that the drinking of heavy drinkers is heavily influenced by TV placements of alcohol, although the modeling assumes that advertising exposure is exogenous: it seems likely that heavy drinkers will select into TV programming that features alcohol so it is unclear what the direction of causality is.

Manning et al. (1995) was the seminal application of QR to this context. This paper used a single cross section of data and a (log) price index based on the weighted average of three drinks (one beer, one whisky and one wine) using average population level shares. Thus, the variation in the price was effectively geographical and it is not possible to separate this from geographic effects so it is unclear the extent to which the estimated elasticity parameters are simply reflecting geographical differences in tastes and incomes. Moreover, there is clearly measurement error in this price and this will imply attenuation - not obviously to the same extent across all quantiles. This paper also considered just one good - units of alcohol. The paper allows for interactions between price and income and nonlinearities in price; and it provides estimates of the probability of drinking (a logit) as well as conditional drinking levels, as well as double log QR estimates conditional on drinking. The paper also finds considerable heterogeneity of the price elasticity: although, unlike Saffer et al. (2012), it finds that the (absolute) price elasticity rises and then falls across quantiles.

New work, by Meng et al (2013) is a lifecycle analysis based on a pseudo panel built by collapsing the data into five year birth cohort groups from just nine years of EFS data. This is an important attempt to unpick cohort and lifecycle effects as well as calendar time effects. However, this lifecycle approach does change the interpretation of the model. In a lifecycle model, estimated price effects reflect the effects of expected variation in prices over time on the planned consumption. They do not estimate the effect of unanticipated changes in prices - for example, price changes arising from changing tax rates. Thus, in itself, this lifecycle approach does not provide estimates that could be informative about tax changes. The construction of pseudo-panels was first done by Browning et al (1985) using a much longer run of Family Expenditure Surveys (FES was the precursor to EFS) which makes very clear the interpretation of such models. Blundell and Walker (1986) also uses FES data, to estimate lifecycle consistent labour supply behaviours, and exploits two stage budgeting.

4 Data

This paper uses data from the Expenditure and Food Survey (EFS) and, its successor, the Living Costs and Food Survey (LCF) and covers the period from 2001, when the diary data first contained quantity information, to the latest survey currently available, 2010. The LCF is the new version of the EFS but the survey has not changed for the purposes of this study. Our focus in on the extent to which the price elasticity varies with the extent of drinking because we are concerned about the relevance of existing UK research for policy in the UK. Thus we implement the QR estimation method, which has been used in the two previous notable US studies, on our UK data that extends the period covered by previous UK studies.

Households are selected if all household members over 15 completed both a face-to-face interview and a two-week expenditure diary. Although individual data is available, our study uses total household expenditure on alcoholic drinks, converted to units of alcohol, because we are concerned about the match between individual consumption and individual expenditure due to intra-household transfers[9]. Since we use household data we need to parameterize the effect of household composition on spending. The extent that we can make inferences about individual consumption from such an aggregate household level model depends on functional form assumptions. Aggregation requires that individual level Engel curves are parallel - something that we cannot test in the absence of reliable individual consumption data[10]. For the moment, we leave the decomposition of household expenditure into individual levels of consumption to subsequent research.

Survey data is likely to exhibit under-recording and we can estimate the extent of this by comparing aggregate expenditure imputed from HMRC recorded excise duties with expenditure obtained from aggregating across the survey data, allowing for the weights in the data. We find that average per adult alcohol consumption in the 2009-10 financial year is 5.9 litres of pure alcohol (590 units) per year in the LCF, compared to the 10.8 litres (1080 units) per adult cleared by HMRC. This is similar to the finding from both the General Lifestyle Survey and Health Survey for England, as reported by Boniface and Shelton (2013). As we noted earlier, if recorded consumption is proportionate to true consumption then estimates obtained from modelling the effect of prices on the recorded data are unbiased. However, the under-recording factor will need to be used to adjust all parameters accordingly and the effects will be the same across quantiles. Thus, this will have no effect on any inferences we make about the extent to which elasticities differ across quantiles.

Quantity information is given in millilitres of drink for every alcoholic drink. As in Purshouse et al (2010), units of alcohol are estimated by multiplying quantity by strength for each drink. We then aggregate across drinks to get a total units of alcohol consumed. The assumptions about strengths are taken

[9]EFS/LCF does not record individual consumption, only individual expenditure and the corresponding quantity purchased. The expenditure data is likely to be contaminated by intra-household transfers - the person who buys is not necessarily the person who drinks. In contrast, the consumption quantity data in GHS is asked at the individual level and seems likely to be a good approximation for true individual consumption. The correlation between individual consumption levels in two-adult households in the GHS data is 0.34 compared with the correlation between individual expenditures of 0.14 in the EFS/LCF data in the same subset of households. The difference between these correlations suggest a great deal of mismatch between individual consumption and individual expenditure

[10]One might be tempted to use the available expenditure data thinking that the resulting measurement error is in the dependent variable and therefore imparts no bias on the estimates. However, we also need individual level income (or, more correctly, individual total expenditure, data). While we might observe income (or total expenditure), it seems likely that expenditure in excess of individual consumption would be reflected in higher total expenditure so any measurement error in the dependent variable would be reflected in the explanatory variable and this would impart some bias. Without a convincing model of the relationship between expenditure and consumption it seems unlikely that one could, at the present time, do better than use the household level data. We can, at least, use the estimated effects of the demographic variables in the model to back out crude predictions of individual consumption.

from their work. If strengths differ across households in ways which are correlated with the explanatory variables then the measurement error in the units of alcohol will be correlated with price. Since the price data is obtained from the household data then the resulting elasticity estimate will be biased.

Purshouse et al (2010) and Collis et al (2010) both generate prices from the diary data, the former using individual-level data and the latter using household-level data. In both cases price is calculated by dividing expenditure by quantity to get a per-unit-price. Using this price as an independent variable, however, risks serious endogeneity since households who purchase a higher quantity of alcohol may pay a lower price. Moreover, forming an explanatory variable by dividing some variable (expenditure) by the dependent variable results in "division bias"[11] since the dependent variable appears on the right hand side of the estimating equation and results in a purely mechanical relationship with the explanatory variable. One option is to proxy price using a price index, which is the method used by Saffer et al. (2012) and Manning et al. (1995) who both use the US ACCRA price index. However, this index has been criticised for not being a very accurate measure of alcohol prices (see Young and Bielinska-Kwapisz, 2003) so inducing measurement error that attenuates the price elasticity estimate. Moreover, the geographical basis for the index implies that it is impossible to extricate the price elasticity in a single cross-section since it will be contaminated by uncontrolled geographical fixed effects. In contrast, the alcohol components of the retail price index provided by the Office for National Statistics (ONS) does not directly have a geographic component to it so the monthly variation in it can be used to identify the price elasticity since the EFS data contains the survey date.

One option is to use the ONS index as an instrumental variable for a price variable computed from the household data to eliminate its potential endogeneity. Another is to use such an index as a proxy variable directly in the estimating equation. A third option is to generate a price index that varies only by month and region but not across households otherwise, by collapsing the raw data to month*region cells, and use this as a proxy or an instrument. There are large regional variations in price, as shown in Table 1. London is clearly more expensive than the North East. It is also interesting to see that the relative prices between on- and off-premise alcohol vary by region, which is likely due to higher fixed costs (such as rent) in regions such as London and the South East. It is clear that there is regional variation in prices, as derived from the microdata itself by dividing expenditure by quantity. While the figures in Table 2 show the price for a unit of alcohol the proportion of each type of drink across households does not vary dramatically across regions suggesting that the price of all drinks vary across regions in line with Table 2.

Table 2 summarises the raw data for drinking households. 19,977 households do not purchase any alcohol. Units refers to units of alcohol. However, the mean values do not show the whole picture. Figure 1 shows a histogram of per-adult units, with reference lines for the median, lower and upper quartiles, 90th and 95th percentiles. The distribution has a long tail, so the distribution has been truncated at 90 units - the 99th percentile of the distribution. A quarter of the sample drink more than 23 units per adult per week. To put this into context, the Royal College of Physicians recommends that men (women) should not drink more than 21 (14) units per week. Multiplying these recommendations by the numbers of adult males and females in the household shows that 35% of households are drinking above these levels. Even this would be a lower bound to the extent of the individual-level problem drinking.

Of course, there are limitations to the dataset. Firstly, the survey is not completely representative of the whole population as it ignores any institutionalised groups such as the military and those in non-residential accommodation - these may be disproportionately heavy drinkers. Secondly, the diary data only covers a period of two weeks, which may mean that observed expenditure is not representative of long-run expenditure. Infrequent purchase may result in a large number not purchasing alcohol in the

[11] See Borjas (1980) in the context of computing the wage rate in a labour supply model by dividing earnings by hours.

Table 1: Average Price per Unit by Region (April 2001 pence)

Region	All Units	On-Premise	Off-Premise
North East	62	105	35
Wales	62	114	37
Yorkshire and the Humber	64	109	38
North West and Merseyside	64	116	37
West Midlands	65	117	37
East Midlands	65	117	37
South West	67	125	38
Scotland	69	135	39
Eastern	71	133	40
South East	75	144	41
Northern Ireland	79	150	41
London	85	159	44

Table 2: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Weekly Alcoholic Units					
Total Household	31.56	34.79	0.25	734	43,036
Total per Adult	17.45	19.48	0.12	734	43,036
On-Premise per Adult	4.69	8.29	0	163.61	43,036
Off-Premise per Adult	12.76	17.88	0	734	43,036
Real Prices per Unit (pence)					
All Alcohol	69.5	59.7	0.64	1635	43,036
On-Premise	128.4	97.4	15.9	2138	29,311
Off-Premise	39.1	17.5	0.1	529	32,957
Household Characteristics					
Adult Males	0.94	0.56	0	8	43,036
Adult Females	0.97	0.51	0	8	43,036
Children	0.59	0.97	0	9	43,036
Age of Oldest Member	51.2	15.6	18	90	43,036
Real Weekly Total Expenditure (£)	372.48	276.81	13.31	7,620	43,036

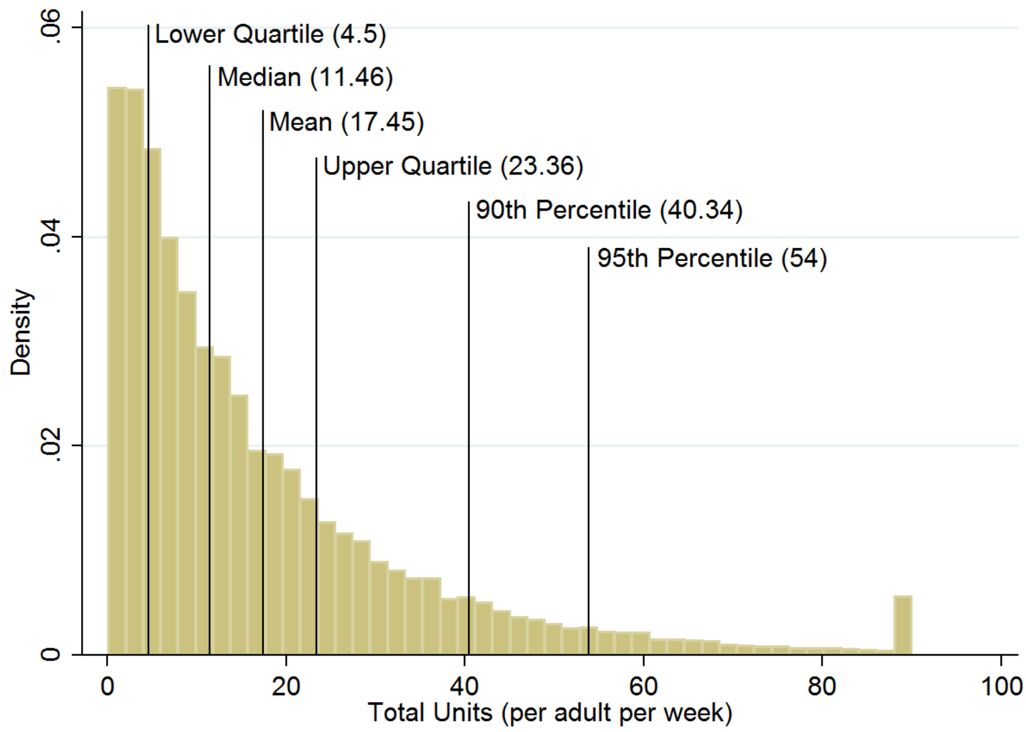

Figure 1: Distribution of Units per Adult per Week

sample period because they are either consuming from stock or because they do not happen to consume within the sample period but nevertheless consume at other times. It is likely that this varies across the year and controlling for month fixed effects captures some of the infrequency. The solution to this problem is to follow Keen (1986), who suggests observed expenditure be replaced by expected expenditure: deflating the positive expenditures by the proportion buying which would leave total expenditure unchanged over the whole sample.

5 Methods

This paper uses quantile regression methods. QR is preferable to subgroup analysis because it does not induce sample selection bias. Running separate regressions for two subsamples essentially truncates the data and biases the estimates. An excellent overview of QR is provided by Koenker and Hallock (2001). The main difference between ordinary least squares and quantile regression is that quantile regression seeks to minimise (weighted) absolute deviations as opposed to the squared deviation. The weight depends on the quantile. This can be written formally, as in Manning et al (1995), as

$$\min_{b_\tau} \left[\sum_{y_i > x_i b_\tau} \tau \, |y_i - x_i b_\tau| \quad + \sum_{y_i < x_i b_\tau} (1-\tau) \, |y_i - x_i b_\tau| \right] \tag{1}$$

where b_τ is the coefficient estimate and τ is the quantile. For example, a median regression would weight both positive and negative deviations equally and set $\tau = 0.5$. Aside from the two alcohol studies mentioned already, QR has been used in only a handful of demand studies: on the demand for tobacco (Goel and Ram, 2004; Chen et al, 2013) and for electricity (Hendricks and Koenker, 1992; Fan and

Hyndman, 2011)[12].

This study uses the double-log model which can be written as

$$lnA_{irt} = \alpha_{irt} + \gamma ln\frac{P_{rt}}{\pi_t} + \lambda ln\frac{EXP_{irt}}{\pi_t} + \beta X_i \quad (2)$$

where A_{irt} is the number of alcoholic units consumed by household i in region r at time t; P_{rt} is the mean price paid over all households within the region r and time period t computed by dividing expenditure by quantity; π_t is the retail price index for all items provided by the Office for National Statistics, EXP_{irt} is total household expenditure, and X is a matrix of demographic, region, and time varying variables. The demographic variables included are the number of children in the household and the age of the oldest household member. Time variables are a linear time trend and monthly dummies. Regional dummies are also included.

Because large households who drink moderately may consume more than a heavy drinking single household, total household units are divided by N^{θ} where N is the number of adults and $\theta \neq 1$ allows for possible non-linear household size effects - for example, individuals who live alone may drink less per capita. When $\frac{A}{N^{\theta}}$ is logged then $\theta logN$ appears as an explanatory variable in the model. We can then use this model to predict per capita consumption within each household for each household's value of N. For the present, we treat men and women the same - more sophisticated analysis could account for the differences between the mix of adult men and women in the household.

We estimate separate models for: all alcohol; alcohol purchased for consumption on licensed premises such as pubs and restaurants (alcohol on); and alcohol purchased away from licensed premises such as shops and supermarkets for consumption elsewhere (alcohol off). There are clear differences in how the prices have evolved over time with the real price of alcohol off falling over the period of our data and real price of alcohol on rising.

6 Results

The results generated by OLS are comparable to the findings in the existing literature, especially the own-price elasticity for all alcohol (-0.586 (standard error 0.039)). QR generates separate estimates of each coefficient for each quantile. We present the three sets of QR price elasticities diagrammatically here in Figures 4, 5 and 6 with their 95% confidence intervals, together with the OLS estimates in each case. We present estimates in the figures of for each vigintile (5% intervals) and present the estimates for the 25th, 50th, 75th, 90th and 95th centiles in the Appendix. The QR results show that 40th percentile has a very similar elasticity as the mean estimate, but that the lighter drinking households are much more responsive to changes in price. The OLS findings are similar to our comparator papers: the finding that demand for on-premise alcohol is more inelastic than the demand for off-premise alcohol is also found Collis et al (2012) and Purshouse et al (2010). Comparing our all-alcohol QR results with those in the US QR literature our findings reflect those in Saffer et al (2012), rather than Manning et al (1995), with significantly smaller than average elasticities above the 60th percentile - but we never find perfectly inelastic demand even for the heaviest drinkers.

[12]It is much more commonly used in the labour economics literature (for example Buchinsky, 1994; Machado and Mata, 2005; Melly, 2005).

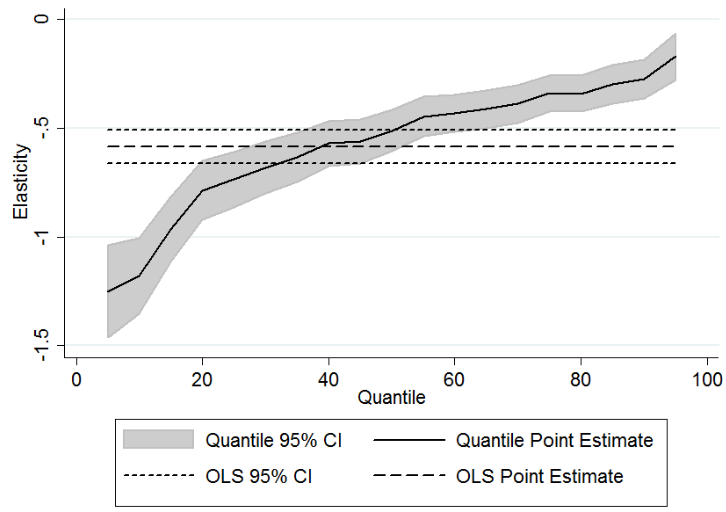

Figure 2: All Alcohol

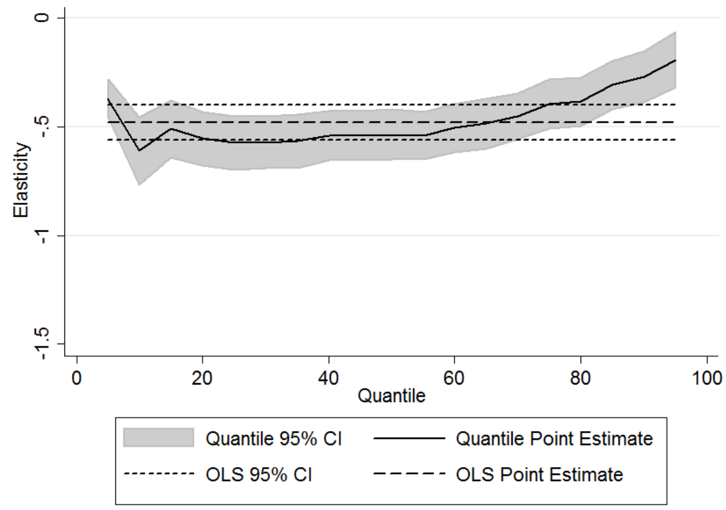

Figure 3: On-Premise

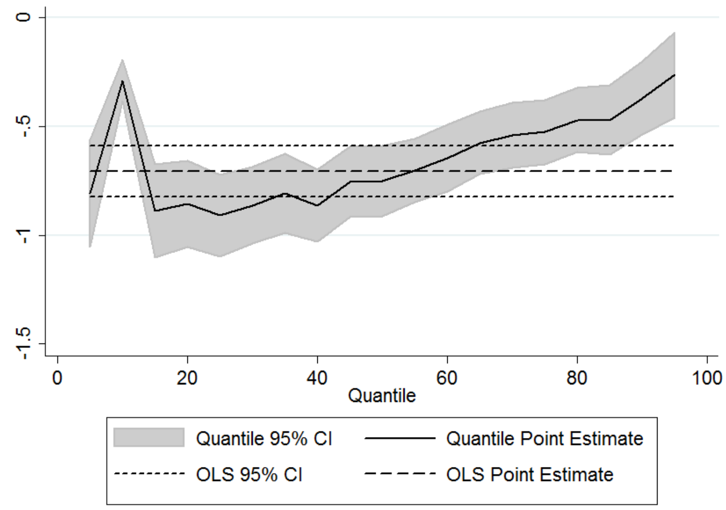

Figure 4: Off-Premise

7 Policy Simulation

The empirical findings suggest that any flat tax levied on alcohol will lead to greater proportional reductions amongst light and moderate drinkers. Attempting to decrease the consumption of the heaviest drinkers through price changes will likely lead to large welfare losses for light and moderate drinkers. The 95th percentile of the drinking distribution consumes 54 units per week per adult compared to the maximum recommended consumption level of 21 units per week (for men), as advised by the Royal College of Physicians. Given that the price elasticity is estimated at -0.17 at this level of consumption, and consumption needs to fall by approximately 61%, price must increase by 359% to bring the 95th percentile in line with the maximum recommended consumption levels. However, this would clearly lead to welfare losses amongst the majority of the population, who are drinking within the recommended limit, as prices rise. The elasticities generated in this paper suggest that a 359% price increase would cause the majority of drinkers to stop consuming altogether. Only the heaviest 15% of drinkers continue to drink. We can estimate the welfare loss caused by a price increase of 359%, as well as the tax revenue raised. Figure 5 shows the effect on consumer surplus of a 359% price increase. The estimates show that welfare loss increases with drinking decile.

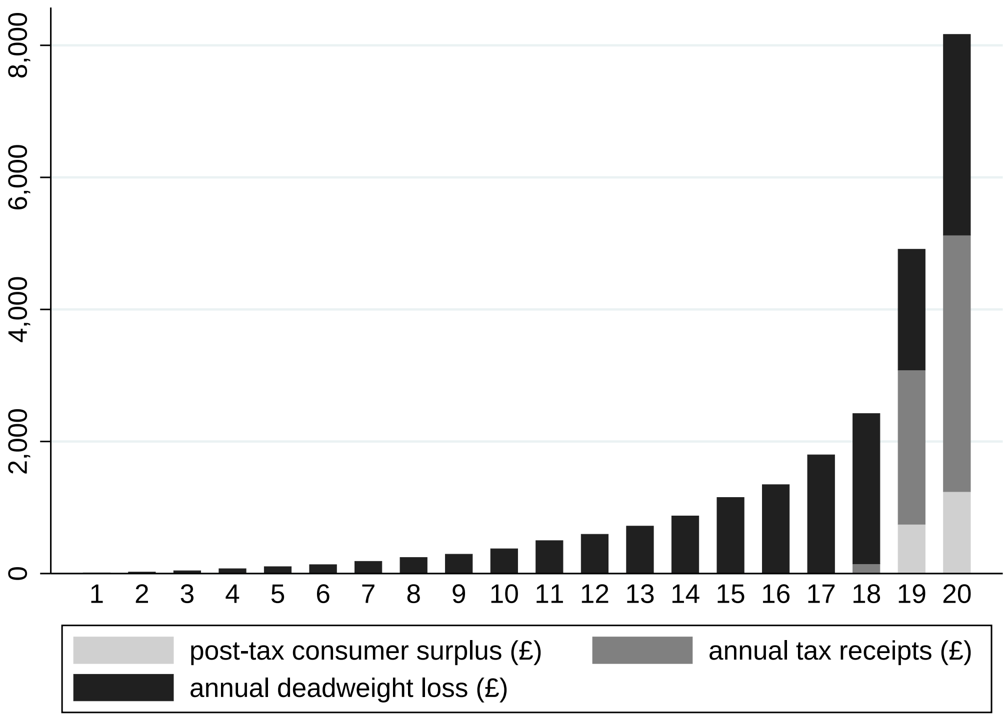

Figure 5: Effects on Consumer Surplus of a 359% Price Increase

Of course, policymakers could target taxes at whichever drink heavy drinkers drink. The data shows that the heaviest drinking decile consume roughly 80% of their units from off-premise drinks, whilst the lightest drinking decile consume roughly 80% of their units from on-premise drinks. The heaviest drinkers also drink a greater proportion of their units in spirits. The on-premise consumption consists mainly of beer, especially amongst the heavier drinkers. Beer accounts for a smaller proportion of off-premise units, with wine accounting for almost half of all off-premise units. Although the lowest off-premise drinking decile consume 7% of their units in ready-to-drink beverages, these account for less than 1% of the heaviest decile's alcohol consumption. Thus it may be possible to exploit the different elasticities associated with each alcoholic produce and location to better target heavy drinkers, although cross substitution effects would need to be taken in account. We intend to pursue this in further research.

The heaviest off-premise drinking decile may also be amongst the heaviest on-premise drinkers. Although 40% of the heaviest off-premise drinkers do not purchase on-premise alcohol, 15% are amongst the heaviest on-premise drinking quintile. It is also worth noting that the heaviest off-premise drinkers pay less per off-premise unit than lighter off-premise drinkers. However, half of the lightest off-premise drinking decile do pay less than 50 pence per unit, and a quarter pay less than 39 pence per unit. A minimum unit price would clearly affect light drinkers as well as heavy drinkers. Of course, another way of interpreting the results is to say that price may not have a very large effect on heavy drinkers. Even doubling the average price would only reduce average consumption from 17.5 units to 11.5 units, and over 16% of drinkers would still be drinking above the recommended amounts. Non-price-based alternatives may be more effective in reducing consumption.

8 Concluding Remarks

This paper has used a large, representative UK dataset to show that heavier drinkers are less responsive to changes in price than lighter drinkers. A 10% price increase would reduce consumption amongst the heaviest drinkers by only 1.7%, compared to over 10% for the lightest drinkers. A price increase of 359% would be necessary to bring the 95th percentile of the drinking distribution to within the maximum recommended level of consumption. The demand for off-premise alcohol, which accounts for the majority of alcohol consumption, especially amongst heavier drinkers, is more elastic than on-premise alcohol. Targeting price increases at off-premise alcohol may be most effective but the price rises would still have to be considerable.

References

Blundell, R. and Walker, I. (1986), "A Life-Cycle Consistent Empirical Model of Family Labour Supply using Cross-Section Data", *Review of Economic Studies*, **53**(4), pp.539-58

Blundell, R., Duncan, A., and Pendakur, K. (1998), "Semiparametric Estimation and Consumer Demand", *Journal of Applied Econometrics*, **13**(5), pp.435-61

Boniface, S. and Shelton, N. (2013), " How is Alcohol Consumption Affected if we Account for Under-Reporting? A Hypothetical Scenario", *European Journal of Public Health*, Advanced Access doi:10.1093/eurpub/ckt016

Borjas (1980), " The Relationship between Wages and Weekly Hours of Work: The Role of Division Bias ", *Journal of Human Resources*, **15**(3), pp.409-23

Browning, M., Deaton, A. and Irish, M. (1985), "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle", *Econometrica*, **53**(3), pp.503-43

Buchinsky, M (1994), " Changes in the U.S. Wage Structure 1963-1987: Application of Quantile Regression ", *Econometrica*, **62**(2), pp.405-58

Chen, C-M., Chang, K-L. and Lin, L. (2013), “ Re-Examining the Price Sensitivity of Demand for Cigarettes with Quantile Regression ”, *Addictive Behaviors*, **38**(12), pp.2801-4

Clancy, G. (2011), “Consumer Trends Quarter 1 2011”, 60, Office for National Statistics

Collis, J., Grayson, A. and Johal, S. (2010), “Econometric Analysis of Alcohol Consumption in the UK”, HMRC Working Paper 10

Cook, P. J. (2008), “A Free Lunch”, *Journal of Drug Policy Analysis*, **1**(1)

Dee, T. S. (1999), “State Alcohol Policies, Teen Drinking and Traffic Fatalities”, *Journal of Public Economics*, **72**(2), pp.289-315

Fan, S. and Hyndman, R. J. (2011), “The Price Elasticity of Electricity Demand in South Australia”, *Energy Policy*, **39**(6), pp.3709-19

Farrell, L. and Walker, I. (1999), “The Welfare Effects of Lotto: Evidence from the UK”, *Journal of Public Economics*, **72**(1), pp.99-120

Fogarty, J. (2010), “The Demand for Beer, Wine and Spirits: A Survey of the Literature”, *Journal of Economic Surveys*, **24**(3), pp.428-78

Gallet, C. A. (2007), “The Demand for Alcohol: A Meta-Analysis of Elasticities”, *The Australian Journal of Agricultural and Resource Economics*, **51**(2), pp.121-35

Goel, R. K. and Ram, R. (2004), “Quantile Regression Estimates of Cigarette Demand Elasticities for the United States”, *Journal of Economics and Finance*, **28**(3), pp.413-21

Hendricks, W. and Koenker, R. (1992), “Hierarchical Spline Models for Conditional Quantiles and the Demand for Electricity”, *Journal of the American Statistical Association*, **87**(417), pp.58-68

Home Office (2012), The GovernmentâĂŹs Alcohol Strategy, London: HMSO.

HMRC (2012), Alcohol Factsheet

Jones, A. M. (1989), “A Double-Hurdle Model of Cigarette Consumption”, *Journal of Applied Econometrics*, **4**(1), pp.23-39

Jones, A. M. and Yen, S. T. (2000), “A Box-Cox Double-Hurdle Model”, *The Manchester School*, **68**(2), pp.203âĂŞ221

Keen, M. (1986), “Zero Expenditures and the Estimation of Engel Curves”, *Journal of Applied Econometrics*, **1**(3), pp.277-86

Kenkel, D. S. (1996), “New Estimates of the Optimal Tax on Alcohol”, *Economic Inquiry*, *34*(2), pp.296-319

Koenker, R. and Hallock, K. F. (2001), “Quantile Regression”, *Journal of Economic Perspectives*, **15**(4), pp.143-56

Leon, D. A., and McCambridge, J. (2006), “Liver Cirrhosis Mortality Rates in Britain from 1950 to 2002: An Analysis of Routine Data”, *The Lancet*, **367**(9504), pp.52-6.

Machado, J. A. F. and Mata, J. (2005), “Counterfactual Decomposition of Changes in Wage Distributions using Quantile Regression”, *Journal of Applied Econometrics*, **20**(4), pp.1099-1255

Manning, W. G., Blumberg, L. and Moulton, L. H. (1995), “The Demand for Alcohol: The Differential Response to Price”, *Journal of Health Economics*, **14**, pp. 123-48

Melly, B. (2005), “Decomposition of Differences in Distribution using Quantile Regression”, *Labour Economics*, **12**(4), pp.577-90

Muellbauer, J. (1976), “Community preferences and the representative consumer”, *Econometrica*, **44**(5), pp.979-99

Purshouse, R. C., Meier, P. S., Brennan, A., Taylor, K. B., and Rafia, R. (2010), “Estimated Effect of Alcohol Pricing Policies on Health and Health Economic Outcomes in England: An Epidemiological Model”, *The Lancet*, **375**(9723), pp.1355-64

Saffer, H., Dave, D. and Grossman, M. (2012), “Behavioral Economics and the Demand for Alcohol: Results from the NLSY97”, *NBER Working Paper 18180*

Symons, E. and Walker, I. (1989), “The Revenue and Welfare Effects of Fiscal Harmonization for the UK”, *Oxford Review of Economic Policy*, **5**(2), pp.61-75

Wagenaar, A. C., Salois, M. J. and Komro, K. A. (2009), “Effects of Beverage Alcohol Price and Tax Levels on Drinking: A Meta-Analysis of 1003 Estimates from 112 Studies”, *Addiction*, **104**(2), pp.179-90

WHO (2011) *Global Status Report on Alcohol and Health*

Young, D. J. and Bielinska-Kwapisz, A. (2003), “Alcohol Consumption, Beverage Prices and Measurement Error”, *Journal of Studies on Alcohol and Drugs*, **64** (2) pp.235-8

Table 3: All Alcohol

		Quantile				
	OLS	25	50	75	90	95
Log Price	-0.586	-0.735	-0.511	-0.340	-0.274	-0.170
	(0.039)***	(0.065)***	(0.049)***	(0.042)***	(0.046)***	(0.054)***
Log Expenditure	0.372	0.375	0.417	0.414	0.381	0.363
	(0.010)***	(0.017)***	(0.013)***	(0.011)***	(0.012)***	(0.014)***
Time Trend	-0.001	-0.002	-0.001	-0.001	-0.001	-0.000
	(0.000)***	(0.000)***	(0.000)***	(0.000)***	(0.000)***	(0.000)
Log Adults	0.347	0.403	0.351	0.286	0.247	0.225
	(0.017)***	(0.028)***	(0.021)***	(0.018)***	(0.020)***	(0.024)***
Children	-0.065	-0.086	-0.069	-0.061	-0.052	-0.049
	(0.006)***	(0.010)***	(0.008)***	(0.007)***	(0.007)***	(0.009)***
Age Oldest	-0.002	-0.004	-0.001	0.001	0.003	0.003
	(0.000)***	(0.001)***	(0.000)	(0.000)***	(0.000)***	(0.001)***
N	43,036	43,036	43,036	43,036	43,036	43,036

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Monthly and regional dummies omitted from output.

Table 4: On-Premise Alcohol

		Quantile				
	OLS	25	50	75	90	95
Log Price On	-0.480	-0.574	-0.535	-0.395	-0.269	-0.192
	(0.042)***	(0.063)***	(0.058)***	(0.058)***	(0.060)***	(0.065)***
Log Price Off	0.023	0.061	0.084	0.003	-0.077	-0.166
	(0.071)	(0.108)	(0.099)	(0.100)	(0.103)	(0.111)
Log Expenditure	0.141	0.153	0.139	0.166	0.153	0.197
	(0.013)***	(0.020)***	(0.018)***	(0.018)***	(0.019)***	(0.020)***
Time Trend	-0.002	-0.002	-0.002	-0.003	-0.003	-0.003
	(0.000)***	(0.000)***	(0.000)***	(0.000)***	(0.000)***	(0.000)***
Log Adults	0.569	0.636	0.718	0.586	0.421	0.318
	(0.020)***	(0.031)***	(0.028)***	(0.029)***	(0.030)***	(0.032)***
Children	-0.150	-0.152	-0.168	-0.163	-0.151	-0.148
	(0.008)***	(0.012)***	(0.011)***	(0.011)***	(0.011)***	(0.012)***
Age Oldest	-0.011	-0.014	-0.014	-0.010	-0.006	-0.004
	(0.000)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***
N	29,311	29,311	29,311	29,311	29,311	29,311

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Monthly and regional dummies omitted from output

Table 5: Off-Premise Alcohol

	OLS	Quantile 25	50	75	90	95
Log Price On	0.079	0.106	0.098	0.027	0.016	0.044
	(0.034)**	(0.055)*	(0.047)**	(0.043)	(0.049)	(0.057)
Log Price Off	-0.707	-0.910	-0.750	-0.526	-0.369	-0.263
	(0.060)***	(0.096)***	(0.082)***	(0.076)***	(0.086)***	(0.100)***
Log Expenditure	0.329	0.330	0.351	0.357	0.326	0.316
	(0.010)***	(0.016)***	(0.014)***	(0.013)***	(0.015)***	(0.017)***
Time Trend	-0.000	-0.000	-0.000	-0.000	-0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Adults	0.138	0.127	0.178	0.194	0.185	0.166
	(0.017)***	(0.027)***	(0.023)***	(0.022)***	(0.025)***	(0.029)***
Children	-0.018	-0.025	-0.017	-0.014	-0.020	-0.023
	(0.006)***	(0.010)**	(0.008)**	(0.008)*	(0.009)**	(0.010)**
Age Oldest	0.005	0.004	0.005	0.006	0.006	0.006
	(0.000)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***
N	32,957	32,957	32,957	32,957	32,957	32,957

Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Monthy and regional dummies omitted from output