

The distributional and nutritional impacts and mitigation potential of emission-based food taxes in the UK

Article

Accepted Version

Kehlbacher, Ariane, Tiffin, Richard, Briggs, Adam, Berners-Lee, Mike and Scarborough, Peter (2016) The distributional and nutritional impacts and mitigation potential of emission-based food taxes in the UK. *Climatic Change*, 137 (1-2). pp. 121-141. ISSN 0165-0009 doi: <https://doi.org/10.1007/s10584-016-1673-6> Available at <https://centaur.reading.ac.uk/63881/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1007/s10584-016-1673-6>

Publisher: Springer

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Abstract

Agriculture and food production are responsible for a substantial proportion of greenhouse gas emissions. An emission based food tax has been proposed as one option to reduce food related emissions. This study introduces a method to measure the impacts of emission based food taxes at a household level which involves the use of data augmentation to account for the fact that the data record purchases and not consumption. The method is applied to determine the distributional and nutritional impacts of an emission based food tax across socio-economic classes in the UK. We find that a tax of £2.841/tCO₂e on all foods would reduce food related emissions by 6.3% and a tax on foods with above average levels of emissions would reduce emissions by 4.3%. The tax burden falls disproportionately on households in the lowest socio-economic class because they tend to spend a larger proportion of their food expenditure on emission intensive foods and because they buy cheaper products and therefore experience relatively larger price increases.

Keywords: food related emission, emission based food tax, Quadratic Almost Ideal Demand System, socio-economic class, data augmentation

1 Introduction

Agriculture and food production including transport, processing, packaging, marketing, sales, purchasing as well as cooking of food are responsible for substantial emissions of greenhouse gases (GHGs). Emissions may occur directly, for example carbon dioxide emissions from fossil fuel use on the farm or in the supply chain, nitrous oxide emissions resulting from fertiliser application, or methane emissions from animals; or indirectly as a result of land use change. In the UK, emissions from food consumption are estimated to contribute 27% of total GHG (167 Mt carbon dioxide equivalent (CO₂e)) emissions (Berners-Lee et al., 2012). According to Gilbert (2012) agriculture is responsible for up to 86% (12,000 mt CO₂e) of all food-related anthropogenic greenhouse-gas emissions, followed by fertilizer production (575 mt CO₂e) and refrigeration (490 mt CO₂e). Several studies find that changes in food consumption behaviour, in particular reduced consumption of meat and dairy foods, can be effective in reducing emissions (CCC, 2010; Edjabou and Smed, 2013; Weber and Matthews, 2008; Stehfest et al., 2009; Garnett, 2011; Vieux et al., 2012; Scarborough et al., 2014) .

Taxing agricultural emissions is one way of including climate change related costs in the market prices of GHG based agricultural products thereby reducing their consumption to socially optimal levels. Emission taxes, however, can lead to emission leakage and according to Schmutzler and Goulder (1997) they are not appropriate where monitoring costs are high, where there are limited options for emission reduction other than by reducing output, or where outputs are highly substitutable. According to Wirsenius et al. (2010) food production fulfils these conditions. In particular in the primary sector there are limited options for adapting the production system to reduce emissions. For example, a change in the ruminant diet, which it would be difficult to foresee more than a small proportion of farmers adopting, would achieve a reduction in methane emissions of only 10% (Hammond et al 2014). The economic argument for food taxes is that the prices paid by consumers do not represent the true social costs of food production. Consequently, consumption of emission intensive foods is higher than is socially optimal. Recent studies investigating the mitigation potential of emission based food taxes conclude that they could reduce emissions, may have beneficial health effects, and are cost-effective. Wirsenius et al. (2010) simulate the impact of an emission based food tax of €60 per tonne CO₂e on ruminant meat, poultry, pork, eggs and dairy products in the European Union (EU). They predict a 7% reduction in food related emissions and conclude that emission based food taxes could change the average diet and could be a cost effective policy for agricultural emission mitigation. In the simulation of the tax induced consumption changes the authors use price elasticities obtained from other studies which they adjust for their purpose, however, some cross price elasticities for example between fish and animal food are assumed to be zero. To compute food related emissions they use FAOSTAT food balance sheets. Briggs et al. (2013) examine the impact of emission based food taxes on health in the UK and conclude that they could significantly reduce food related GHG emissions whilst at the same time improving population health and generating substantial tax revenue. The present study uses a similar approach to Briggs et al. (2013) which involves choosing a tax rate based on information about marginal abatement costs, estimating a demand system that is more detailed than the one estimated by Wirsenius et al. (2010) including for example fish and beverages, and estimating price elasticities using household data.

The present study differs from Wirsenius et al. (2010) and Briggs et al. (2013) in two respects. First, the previous studies provide emission estimates based on changes in the average diet of the entire population thereby assuming that all households respond to tax induced price changes in the same manner. In practice, households are likely to respond differently to the tax depending on how much of a food with a given tax rate they buy, whether they buy cheap or expensive products in the taxed food categories, and the resources available to them to compensate for the food price increases. Our approach is therefore to compute the impacts of the tax at the level of the individual household in our sample. As well as allowing for differences in responses across households we are able to use this approach to compute the distributional impacts of the policy. Like carbon taxes on energy (Wier et al., 2005), emission based food taxes may be regressive as several studies find an association between socio-economic status and diet (e.g. Billson et al., 1999; De Irala-Estevez et al., 2000; Turrell and Kavanagh, 2006; Tiffin and Arnoult, 2010). To analyse the distributional impacts we must account for the fact that our data record purchases, not consumption, over a two week period. An important reason why these may differ is where the household holds food stocks. Where stocks increase, purchases over-estimate consumption and therefore emissions and the reverse is true when consumption occurs from stocks. Our second contribution is therefore to show how a by-product of our estimation algorithm can be used to correct the bias that would arise if purchases were used to calculate emissions at the household level.

2 Data

2.1 Consumption Data

We use data from the 2011 UK Living Costs and Food survey (LCF, 2011). The data are a cross-section of 5,691 households which are chosen using multi-stage stratified random sampling to ensure that the sample is representative of the UK population. Data on food expenditures and quantities purchased for consumption at home are collected from each household over a two week period and these periods are distributed across the year for different households. Because we are interested in the distributional impacts as well as how the tax affects the diets of households with low socio-economic status which tend to consist of less healthy foods (Darmon and Drewnowski, 2008; Harding and Lovenheim, 2014), we group households according to socio-economic class as follows: higher managerial & professional (SEC1), lower managerial & professional (SEC2), intermediate, small employers & own account workers (SEC3), lower supervisory & technical, semi-routine & routine (SEC4), and occupations not stated, not classifiable for other reasons (OTHER). Sample characteristics are reported in Table 5 in the appendix as well as each group's population equivalent as computed by the household specific population weights. OTHER has the highest average pensioner income and number of pensioners per household and it contains mostly older household member who are economically inactive. For our analysis this means that the households in OTHER are heterogeneous in terms of SEC, they are included in the analysis to ensure that the sample is representative.

2.2 Carbon conversion factors

To compute emissions we use conversion factors (Table 1) obtained by using estimates of GHG emissions up to the point of sale at a mid-sized supermarket chain in the northwest of England with 26 stores. Thus, it is assumed that the GHG emissions embodied in this supermarket chain's product range are representative of those sold by all other food retailers in the country. The calculation of embodied GHG emissions follows the principles of the Greenhouse Gas Protocol (World Business Council for Sustainable Development, 2004) and draws upon a range of secondary sources of life cycle analysis studies. The carbon conversion factors include components for farming and manufacturing, transport, packaging, storage and supermarket operations. The latter includes transport to store as well as energy consumption, staff business travel, postage and courier services, waste disposal, paper, printing and other office and marketing consumables. Emissions from land use change are taken into consideration by using carbon conversion factors that factor them in for key foods such as beef. The methodology used for the input–output analysis is described in detail by Berners-Lee et al. (2011) and for the derivation of the carbon conversion factors the reader is referred to Berners-Lee et al. (2012).

The conversion factors vary between groups for some food categories because they are weighted averages of several conversion factors. For example, the conversion factor of the other milk category is a weighted average of factors for yoghurt and fromage frais, ice cream and desserts. The weights are the average expenditure shares of the constituent food categories. The carbon conversion factors are applied to the (latent) quantities to compute food related GHG emissions as described in Section 3.2 and for each household the product of quantity and conversion factor is multiplied by a household specific population size weight to produce a population estimate of food related GHG emissions. The weights ensure that the sample matches the UK population in terms of region, age group and sex.

3 Method

3.1 Demand model estimation

We adopt a hierarchical approach to the estimation of the price elasticities which requires a number of separability assumptions to keep the dimensions of the estimated models manageable. Foods are grouped into nutritionally meaningful categories and according to their embedded emissions (see Figure 1 in appendix). The model at the top level represents the household's decision to allocate overall food expenditure between broad food categories including drinks (see Table 6). Next, seven models represent the household's decision to allocate expenditure between the dairy & eggs through to the alcohol categories. Below this, the starches, fats and beverages categories each have a model explaining the budget allocation between a further disaggregation, for example, into coffee and tea & cocoa drinks. A Quadratic Almost Ideal Demand System (QUAIDS) (Deaton and Muellbauer, 1980; Banks et al., 1997) is estimated for each of the eleven demand systems each of which has m equations. The QUAIDS was chosen because it allows for non-linear Engel curves which means expenditure elasticities can vary with expenditure levels. The estimation method we employ is detailed in Tiffin and Arnoult (2010) and has subsequently been applied by others including Kasteridis et al. (2011). Our accounts for censoring in the expenditure shares (see Online Resource 1) by adapting the model to incorporate infrequency of purchase (Cragg, 1971; Blundell and Meghir, 1987). A system of probit and demand equations is estimated

$$\begin{pmatrix} y^* \\ s^* \end{pmatrix} = \begin{pmatrix} X_1 & 0 \\ 0 & X_2 \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} + \begin{pmatrix} u \\ v \end{pmatrix} \quad (1)$$

where y^* is a vector of latent probit variables, s^* is a vector of latent expenditure shares (see also section 3.2),

X_1 is a vector of constants, X_2 is a matrix of prices and household expenditure. Expenditure shares are expressed as follows

$$s_{ih}^* = \alpha_i + \sum_{j=1}^m \gamma_{ij} \ln(p_{jh}) + \beta_i \left(\frac{e_h}{a(p)} \right) + \frac{\lambda}{b(p)} \left\{ \ln \left(\frac{e_h}{a(p)} \right) \right\}^2, i = 1, \dots, m \quad (2)$$

where

$$a(p) = \exp \left(\alpha_0 + \sum_{i=1}^m \alpha_i \ln(p_{ih}) + \frac{1}{2} \sum_{i=1}^m \sum_{j=1}^m \gamma_{ij} \ln(p_{ih}) \ln(p_{jh}) \right), b(p) = \prod_{i=1}^m p_i^{\beta_i} \quad (3)$$

p_{ih} is the price faced by household h for food category i , e_h is food expenditure of household h , and α and β are the coefficients to be estimated. Adding up, symmetry, homogeneity and concavity are imposed in the estimation. Because the covariance matrix is singular we drop one of the share equations, obtaining estimates that are invariant to which equation is dropped (Barten 1969). Food categories are aggregated using the EKS quantity index (Elteto and Koves, 1964; Szulc, 1964), which is a multi-lateral version of the superlative Fisher Ideal index, which is used to compute the implicit price index. A superlative index (Diewert, 1976) offers some mitigation towards the concerns over the potential endogeneity of prices.

We use Markov Chain Monte Carlo (MCMC) methods to sample from the posterior distribution of the parameters. MCMC methods are commonly used to sequentially draw values from an approximation to the posterior density which is proportional to the product of the likelihood function and prior distributions of the parameters. Because the distribution of a sampled draw depends only on the last value drawn, the draws form a Markov Chain and under the right conditions these draws converge to represent draws from the full posterior. We use Gibbs sampling to draw subsets of parameters using their conditional posterior distributions (Gelman et al., 2014).

The data is censored and we therefore use the infrequency of purchase model (Cragg, 1971; Blundell and Meghir, 1987) which separates the discrete buying decision from the continuous decision on how much to buy. Accordingly, two types of latency arise. Both stem from the fact that observed purchases differ from consumption when households consume from stocks. The first type of latency arises where a purchase is observed and it must be deflated to yield an estimate of the latent quantity consumed:

$$q_{ih}^* = q_{ih} \Phi_{ih} \quad (4)$$

where q_{ih} is the observed purchase and $\Phi_{ih} = P(q_{ih} > 0) = P(y_{ih}^* > 0) = \Phi(x_{1h} \beta_{1i})$ is the probability that a purchase is made. Using the estimate of latent consumption the latent share is obtained:

$$s_{ih}^* = \frac{p_{ih} q_{ih}^*}{\sum_{i \in C} p_{ih} q_{ih}^*} \quad (5)$$

where C is the set of observations that are not censored. Where food is consumed from stocks but no purchase is observed the second type of latency occurs. We treat this as an incomplete data problem and replace the censored values using data augmentation (Tanner and Wong, 1987). Thus, where household h makes no purchase of food category i the observed zero is replaced with a latent share s_{ih}^* drawn from the conditional distribution:

$$s_{ih}^* \mid \beta_1, \beta_2, \Sigma, D \sim N(\mu_{ih}, V_i) \quad (6)$$

where μ_{ih} and V_i are the conditional mean and conditional variance respectively and D is observed and latent data. Observed shares are scaled using an adapted Wales and Woodland (1983) procedure as described in Tiffin and Arnoult (2010) to ensure that adding up holds.

We draw 12,000 MCMC samples and discard the first 2,000 draws to ensure convergence. The remaining 10,000 draws taken from the posterior provide the basis for inference. As convergence diagnostics we use trace plots and the Geweke test (Geweke, 1992) which compares the means from the first and second half of the Markov chain.

Unconditional price elasticities (Online Resources 2) are computed for each group as the mean of the household specific elasticities within the group (Online resource 3). These assume that a price change of one of the food categories changes food expenditure available to all other food groups. The elasticities are used to compute the changes in food consumption. Changes in nutrient intake are computed using the corresponding nutrient conversion factors provided by the LCF survey.

3.2 Computing food related emissions

Previous studies have computed food related emissions based on food purchases. This is problematic when the frequency of purchase differs across socio-economic groups. Consider two households that eat steak once a week and which are representative of two different socio-economic groups. One household shops for steak once every three weeks and the other shops once a month. The level of emissions that we compute depends on which households happened to shop in the two week period in which the data is collected. The latent quantities which are computed in the course of estimating the model offer a solution to the problem. Because these have been adjusted taking into account the probability that a purchase occurred in the survey period, a true representation of emissions can be obtained. For cases where no purchase is made and the latent share is obtained using equation 6, the latent share may be negative. Where this occurs we assume that consumption is zero and therefore no emissions arise. Overall, emissions g_{ih} are:

$$g_{ih} = \begin{cases} q_{ih}^{**} F_i & \text{if } s_{ih}^* > 0 \\ 0 & \text{if } s_{ih}^* \leq 0 \end{cases} \quad (7)$$

where

$$q_{ih}^{**} = \frac{s_{ih}^* \sum_{l \in C} p_{lh} q_{lh}^*}{p_{ih}} \quad (8)$$

and F_i is the carbon conversion factor of food category i .

3.3 Determining the tax rate and proportional price increases

Table 2 shows the tax rate for each food category and the resultant proportional price changes for Scenario A which imposes a tax on all foods according to their emission content. Because the administrative costs of taxing foods with low emissions would be disproportionately high, we also simulate a tax only on foods with above average levels of emissions (Scenario B).

The tax rate is computed using the agriculture marginal abatement cost curve (MACC) provided by the Department for Environment, Food and Rural Affairs (Moran et al., 2008) which suggests that investment of £24.10/tCO₂e (£28.41/tCO₂e, 2011 prices) can reduce UK agricultural GHG emissions by 16.2%, and from the Stern Review's (Stern, 2007) social cost of carbon which is calculated as £22-£26/tCO₂e (2011 prices) emitted to maintain global atmospheric concentration of carbon dioxide equivalents at 450-550 parts per million. The marginal external cost is assumed to be £2.841/tCO₂e per 100g of product (Moran et al., 2008). The next most cost-effective abatement strategy suggested by the MACC (Moran et al., 2008) is £205.40/tCO₂e at 2011 prices and is prohibitive. The tax rate chosen is therefore a compromise between the practicality of abatement and societal costs. To obtain the tax rate we multiply the carbon conversion factor by the marginal external cost of emissions of £2.841/tCO₂e. The proportional price changes are obtained by comparing the average price per 100 g of food with the new price after the tax has been added. This is done for each food category and also for each group because average prices differ between groups. Households that buy cheaper products of a given food category and therefore have lower average expenditure per food unit, experience larger proportional price increases, and vice versa. For example, SEC4 tends to have the smallest unit values for many food categories and therefore experiences the largest proportional price increases in most cases. By contrast, SEC1 has the smallest proportional price increases as people working in higher managerial & professional functions are expected to have a higher income (see Table 5) and spend more on food (see Table 6) which is likely to translate into them buying more expensive products.

The impact of the tax also depends on the proportion of food expenditure that a specific socio-economic group devotes to food categories that are emission intensive and therefore attract a higher tax rate. The expenditure shares show that the OTHER group tends to have the largest expenditure shares for emission intensive food categories such as milk, pork, beef, lamb, other meat and therefore its tax burden will be higher. By contrast, SEC1 has larger expenditure shares for less emission intensive foods such as fresh fruit, fresh vegetables, and fruit juice, and both SEC1 and SEC2 have smaller expenditure shares for meat and dairy foods. Thus, higher SECs will be less affected by the tax because they buy relatively more foods with low emission intensity while SEC4 will be most affected because it consumes relatively more emission intensive foods and experiences larger proportional price increases.

4 Results

4.1 Changes in food consumption and nutrients

The estimated change in the consumption of a given food category depends on the own and cross price elasticities, (Online Resources 2), the magnitude of the proportional price increases, and the expenditure share. For example, in Scenario A (see Table 7) eggs are taxed and all groups reduce their consumption by 4-7 per cent. In Scenario B eggs are not taxed but consumption is also reduced by 3-7 per cent because of a strong complementarity between eggs and cheese, which is taxed in both scenarios. In the case of vegetable fats, consumption increases in Scenario A for all groups except SEC4 despite being taxed. This is because of the strong substitution effect between vegetable fat and the bread etc. and cakes etc. categories. The changes in fruit and vegetable purchases are also of note. Their consumption falls in both scenarios although they are not taxed in Scenario B. This is because of complementary relationships that exist between the fruit and vegetable categories and other food categories.

In both tax scenarios, SEC4 and OTHER decrease their food consumption the most. SEC4 and OTHER differ in that OTHER includes mostly pensioners who have diverse socio-economic backgrounds and as a result their preferences and responses to price changes are more varied. Although SEC4 and OTHER both have low incomes, the differences in price elasticity between the two groups imply that they are very different in their food preferences and in their food shopping behaviours and therefore in their responses to the tax.

In addition to the substitution effects arising through differences in the relative prices of food, the imposition of a tax brings about a general price rise and hence a reduction in real expenditure on food. Consequently the tax also causes an overall reduction in food consumption. In Scenario A (see Table 7), SEC4 reduces its food consumption in most instances by the largest amount because it has the largest price elasticities and often experiences the largest proportional price increases. SEC1 and SEC2 tend to decrease consumption to lesser extents. In Scenario B (see Table 8) SEC4 also strongly decreases consumption as well as SEC2 which decreases coffee drinks, canned and fresh veg and lamb consumption because of its having complementary relationships between these food categories and the taxed food categories. A detailed discussion of the consumption changes is given in Online Resource 5.

For both tax scenarios, changes in average daily intake of selected nutrients are reported in Table 9 (absolute quantities in Online Resource 4). The impact of the tax on diet quality is mixed. Households reduce consumption of energy, fat, sugar and salt but they also reduce consumption of fibre as a result of the reduction in fruit and vegetables. The proportionate reductions in fruit and vegetable intake are slightly below those in energy, fat, sugar and salt intake which could be interpreted as a marginal improvement in the balance of the diet.

4.2 Greenhouse gas emissions and tax revenue

Using the carbon conversion factors, we compute per capita household emissions for the average diet. SEC1 and SEC3 have the most emission intensive diets with 4.17 gCO_{2e}/kcal and 4.15 gCO_{2e}/kcal, respectively. These are followed by SEC2 (4.12 gCO_{2e}/kcal) and SEC4 (3.91 gCO_{2e}/kcal) while OTHER has the least emission

intensive diet (3.91 kgCO₂e/kcal). The apparent contradiction with the fact our earlier statement that OTHER spends the largest share of total expenditure on emission dense food is explained by the fact that this group consumes smaller absolute quantities of these foods.

Table 3 shows that estimated total emissions from food consumed at home are 126 MtCO₂e. This constitutes 22% of total the UK GHG emissions of 577.3 MtCO₂e for 2012 (Webb et al., 2014). In comparison, Audsley et al. (2010) estimate that the supply of food and drink for both at home and away consumption for the UK results in direct emissions of 152 MtCO₂e, and Berners-Lee et al. (2012) estimate the amount of GHGs embodied in the UK's food supply to account for 27% of total emissions.

Taxing all foods (Scenario A) reduces food related emissions to 118 MtCO₂e (118 to 120 MtCO₂e 95% credible intervals) which constitutes an average decrease of 6.3% (8 ktCO₂e). The contribution by food and agriculture to UK total emission is reduced from 22% to 20.5% using emission estimates by Webb et al. (2014) for the year 2012. An above average contribution to this reduction is made by SEC4 which reduces emissions by 7.3%, followed by OTHER which reduces emissions by 6.5%. The lowest reduction in emissions is achieved by SEC1 and SEC2 (-5.4% and -5.5% respectively). Annual tax revenue from an emission based food tax on all foods is estimated be £6.06 billion (£5.97 to £6.40 billion, 95% credible intervals).

After introducing a tax on foods with above average levels of emissions (Scenario B), emissions are reduced to 121 MtCO₂e (120 to 122 MtCO₂e, 95% credible intervals) which constitutes a decrease by on average of 4.3% (5.5 ktCO₂e). Using emission estimates by Webb et al. (2014), this means that the contribution to UK total emission decreases from 22% to 20.1%. As in Scenario A, SEC4 reduces emissions the most but its reduction is closely followed by OTHER and SEC3 with all three household groups reducing emissions by around 4%. By contrast, SEC1 and SEC2 reduce emission by around 3% which is similar to the reduction achieved in Scenario A. The revenue from a tax on foods with above average emissions levels is estimated to be £3.61 billion (£3.55 to £3.77 billion, 95% credible intervals).

5.1 Discussion

In this study we use conversion factors which are UK specific and closely match the food categories that are estimated in the demand systems. Our study differs methodologically in that it uses household data and latent quantities to compute emissions where no purchases are made, thereby giving an accurate measure of baseline emissions which is larger than it would be if zero purchases were treated as zero consumption. The study also differs in that it estimates the QUAIDS and it accounts for censoring in the estimation procedure. We find that an emission based food tax in the UK on all foods reduces GHG emissions from food consumed at home by on average 6.3% (-8.023 MtCO₂e) and a tax on foods with above average levels of emission alone would reduce emission on average by 4.3% (-5.492 MtCO₂e).

In comparison, Briggs et al. (2013) estimate an emission based food tax of £2.719/tCO₂e on foods with above average emissions and find a reduction potential of 7.5%. Our lower estimates are largely due to our conversion factors being smaller¹. These differences influence the results in three ways: i) because the emission tax is a levy per unit of food, smaller carbon conversion factors imply lower tax rates; ii) smaller carbon conversion factors

¹ For example, in case of beef, lamb, other meat and animal fats, Briggs et al. use factors of 68.8 kgCO₂e, 35.9 kgCO₂e and 35.6 kgCO₂e respectively.

imply less emissions computed from the total quantity of food consumed; and iii) for Scenario B different conversion factors imply that different types of foods are classed as having above average levels of emissions. Differences between conversion factors are difficult to explain because their derivation is not explained in detail, but the emissions factors used in this study were selected on the basis of reliability and transparency of their derivation meaning they can be traced back to peer reviewed papers which include detailed methodologies and transparent descriptions of their derivation (see Hoolohan and Berners-Lee, 2012 for details). Compared to Wirsenius et al. (2010) who report a 7% reduction of GHG emissions in the EU by taxing animal foods, our estimates appear to be considerably larger given that our tax rate is about 15 times lower than theirs. Our conversion factors capture emissions not only from agriculture but also from manufacturing, transport, packaging, storage and supermarket operations. More importantly the difference in predicted emission reductions is due to differences in the estimated price elasticities with our price elasticities being generally larger which indicates that households are more responsive to the tax than previously suggested. The difference in price elasticities could be explained by differences in aggregation schemes. Whilst Wirsenius et al. estimate elasticities for 7 categories excluding fish, fruits and drinks, our demand system has 29 categories including non-alcoholic drinks, alcohol, fish and fruits. Higher disaggregation leads to larger elasticities because there are fewer opportunities for substitution within a given food category. Also, there are differences in the time periods covered and model specifications.

In both tax scenarios the tax burden disproportionately falls on households in the lower SECs which makes emission based food taxes regressive. SEC4 is most affected followed by the socio-economically heterogeneous group OTHER. SEC4 is more affected because its households tend to buy cheaper products and therefore experience larger proportional price increases. At the same time, SEC4 has the lowest average weekly income and therefore has fewer resources to absorb higher food prices. OTHER is more affected because its households have the largest expenditure shares for emission intensive food categories such as meat and dairy and also starches. By contrast, households in higher SECs tend to have diets with less meat and dairy foods and more fruit and vegetables and therefore they are less affected by the tax. These households also buy more expensive products and therefore experience smaller proportional price increases. SEC1 has the highest household income and experiences the smallest proportional price increases for most food categories which means it is least affected by the tax, followed by SEC2. This is in line with research by Feng et al. (2010) who investigate the distributional effects of climate change taxes on households belonging to different income and lifestyle groups and find them to be regressive.

The impact of the GHG tax on dietary quality is mixed. The positive health impact from reducing the intake of fat, sugar and salt maybe offset by a reduction in fibre, fruits and vegetables in particular in case of SEC4. The tax is not efficient in the sense that it does not incentivise households with emission intensive diets to reduce their emissions relatively more because income and cross price effects between all food categories confound its impact.

5.2 Limitations

Some caution is warranted when interpreting these findings. First, due to data limitations this study estimates unconditional elasticities only for food, thereby assuming that total expenditure on food stays constant. In practice, households may compensate for food price increases by shifting some of their non-food expenditure towards food. Our estimated reductions in food consumption and GHG emissions therefore may constitute upper boundaries. Second, we assume that the tax is passed on in its entirety to consumers. Retailers, however, may spread the increased cost of the taxed products over other products thereby attenuating the effect of the tax; or retailers may use their market power to pressure suppliers to reduce costs in their production process such as costs for labour or animal welfare. For example, in the case of a tax on sugar-sweetened beverages, Cawley and Frisvold (2015) find that prices rose by less than half of the amount of the tax. Third, consumers may respond to the tax by increasingly buying bulk and special offers or cheaper or lower quality products to compensate for higher prices. Our data show that higher SECs buy more expensive food products which means they have some capacity to absorb price rises by switching to cheaper products within the same emission intensive food category. The effect of this kind of inter food category substitution on emissions is uncertain because it depends on the emission intensity of the substitutes. Fourth, the conversion factors are subject to uncertainties as they are derived in relation to a specific case study UK supermarket supply chain the characteristics of which may differ from other UK supermarket supply chains. Fifth, the use of population size weights to scale quantities up for the entire population means that food consumption patterns and responses to price changes by sample households are assumed to be representative of all households in the population who have the same characteristics. Sixth, the response by households to a tax may differ from their response to a price change. Finally, while our method of computing food related emission from latent shares accounts for the uncertainty surrounding the sources of zero purchases in the data, it is computationally intensive and therefore may not be suitable for larger demand systems. To simplify the estimation of the demand system we have not included socio-economic variables which would allow to the probability of purchase to vary within each SEC.

5.3 Policy implications

While emission based food taxes could make an important contribution to the reduction of GHG emissions in the UK, they may be difficult to impose in practice because the tax burden falls disproportionately on households in lower SECs. Unintended health consequences can also arise as the tax fails to take into account that some individuals may benefit from consuming emission intensive foods such as milk. Equally taxing foods according to their emission content can create perverse incentives with energy dense foods such as sweets and soft drinks attracting lower tax rates than more nutritious foods. One might argue that despite the inefficiencies that arise with the tax, the revenue raised could be used to expedite low emission innovations in food production. However, if this is the goal, revenue should be raised in the most efficient and least regressive way. Government should not confound the revenue raising objective of a fiscal instrument with behaviour change. Other ways of reducing food related emissions should also be considered such as the reduction of food waste in food processing, distribution, retail and the home, or the use of carbon labelling or nudges to alter households' food shopping behaviour.

6 Conclusions

This study investigates the GHG emission reductions that would be achieved by introducing an emission based food tax in the UK. To analyse the distributional impacts we use household level purchase data. In the computation of food related emissions we account for the fact that emissions are based on food consumption which we infer from the purchase data by using the outputs from the data augmentation procedure in our estimation algorithm. We find that an emission based food tax would affect households in the lowest SEC the most because they spend a larger proportion of their food expenditure on emission intensive foods such as meat and dairy even though they buy smaller absolute quantities. Households in the lowest SEC also buy cheaper products and therefore experience larger proportional price increases. An emission based food tax is therefore regressive and may require some form of compensatory mechanism to make it socially acceptable. For example, a tax schedule akin to the British Columbia carbon tax which achieves carbon emission reductions by taxing carbon and using the revenue to lower income and corporate taxes (The Economist, 2014) could be implemented. We also find that the impact of the policy on dietary quality is ambiguous. Consumption of undesirable nutrients will fall but healthy components of the diet, in particular fruit and vegetables, also decline. Overall, we conclude that whilst an emissions based food tax could reduce GHG emissions, there are some significant risks associated with such a policy.