

Bardsley, Nicholas; Büchs, Milena; Schnepf, Sylke V.

Working Paper

Something from Nothing: Estimating Consumption Rates Using Propensity Scores, with Application to Emissions Reduction Policies

IZA Discussion Papers, No. 9707

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Bardsley, Nicholas; Büchs, Milena; Schnepf, Sylke V. (2016) : Something from Nothing: Estimating Consumption Rates Using Propensity Scores, with Application to Emissions Reduction Policies, IZA Discussion Papers, No. 9707, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/141466>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 9707

Something from Nothing: Estimating Consumption Rates Using Propensity Scores, with Application to Emissions Reduction Policies

Nicholas Bardsley
Milena Büchs
Sylke V. Schnepf

February 2016

Something from Nothing: Estimating Consumption Rates Using Propensity Scores, with Application to Emissions Reduction Policies

Nicholas Bardsley

University of Reading and Walker Institute for Climate Change Research

Milena Büchs

University of Southampton

Sylke V. Schnepf

European Commission JRC, University of Southampton and IZA

Discussion Paper No. 9707
February 2016

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Something from Nothing: Estimating Consumption Rates Using Propensity Scores, with Application to Emissions Reduction Policies*

Consumption surveys often record zero purchases of a good because of a short observation window. Only mean consumption rates can then be inferred. We show that propensity scores can be used to estimate each unit's consumption rate, revealing the distribution. We demonstrate the method using the UK National Travel Survey, in which c.40% of motorist households purchase no fuel. Estimated consumption rates are plausible judging by households' annual mileages, and highly skewed. We apply the same approach to estimate CO2 emissions and direct outcomes of a carbon cap or tax. Analysis of such policies based solely on means appears to have a negative bias, because of skewness of the underlying distributions. The regressiveness of a simple tax or cap is overstated, and redistributive features of a revenue-neutral policy are understated.

JEL Classification: C13, D04, D12, H23

Keywords: emissions reduction, fuel consumption, infrequent purchase, propensity score matching, surveys

Corresponding author:

Nicholas Bardsley
School of Agriculture, Policy and Development
University of Reading
PO Box 237
Reading RG6 6AR
United Kingdom
E-mail: n.o.bardsley@reading.ac.uk

* The authors would like to acknowledge comments from participants in seminars at the University of Reading, UK, and at the European Commission Joint Research Centre in Ispra, Italy, from Mr. Brian Davey and Dr. Simone Pfuderer. The idea of using propensity scores in this context was suggested by Dr James Brown at S3RI, University of Southampton, UK (personal communication). This research was funded by the ESRC, grant RES-000-22-4083.

NON-TECHNICAL SUMMARY

While there is an overall agreement that climate change needs to be restricted, the best choice of policies for achieving this goal is uncertain. Standard policy proposals include carbon taxes and carbon rationing policies, but it is unknown whether these policies distribute the costs of CO₂ reduction fairly across the population. Carbon-intensive goods are purchased too infrequently for estimating their use by households with different characteristics (like poor or rich). This paper shows how this methodological problem can be overcome, using the example of motor fuels based on UK data.

Based on a consumption sample 40% of motorists purchase no fuel within one week of the survey window. However, these motorists can be matched to households that did buy fuel, based on their likelihood of purchase. This allows us to estimate fuel consumption rates for each household in the sample, revealing the distribution of consumption. The distribution turns out to be highly skewed, with relatively few high-emissions households responsible for a very large share of emissions. The top 20% of motorist households account for around 40% of vehicle fuel emissions, around 11 times the share of the bottom 20%.

Our improved knowledge of the distribution of consumption rates can be used to analyse carbon taxes and rations. Without reconstructing the distribution, analysts previously had to rely on mean estimates. But these are misleading because of the skewness prevalent in consumption data which requires information on the median or middle consumer. This paper shows that by focusing on the median instead of the mean using the methodology proposed revenue-neutral emission reduction policies are much fairer in terms of redistribution than was previously assumed.

1. Introduction

A problem in survey sampling is that events of interest may occur infrequently relative to the data collection period, leading to zero-inflated data. There are many circumstances in which this can arise. For example, if a wildlife survey runs camera traps for a short time, negative results may obtain even where a target species is known to reside (Burton *et al.*, 2015). In the social sciences a key example is purchase infrequency. This occurs when a diary is used to record buying over a relatively short duration, typically 1-2 weeks. Sampled households often record no purchase, even if they are known to consume the good in question. All households consume clothing, but many will not make purchases in a given two weeks, for example. For many questions of social and political interest, there is as we shall explain a consequent paucity of information. An example analysed later in this paper is CO₂ emissions reduction policy, including carbon taxes or caps. This is topical given that the recent COP21 meeting ended with agreed aims to restrict global warming, but without agreement on policies to achieve this.

Consumption surveys aim to measure rates of consumption, which unlike quantities purchased do not depend on the length of the observation window. The same drinking rate, for example, can be expressed as 1 pint per day, 7 pints per week or 365 pints per year. In a well-designed and executed survey, the estimated mean consumption rate is not biased by purchase infrequency. For zeroes will tend to be counterbalanced by positive values that, if interpreted as consumption rates, would be too high.¹ The problem is that the data are uninformative about a given household's consumption rate, and therefore about any other statistic than the mean. Concerning emissions reduction policies, one would like to judge their potential regressiveness and likely policy acceptance. But this requires knowledge of key carbon-intensive purchases, including flights and motor fuels, which are relatively infrequent and therefore of unknown distribution.

¹ For example, if each household used 1, 2l bottle of milk every 2 weeks, an ideal survey with a 1 week diary would be expected to record 50% of units with no purchase and 50% purchasing 2l. The expected mean rate of 1l per week is correct, but no entry would record 1l.

In this paper we apply Propensity Score Matching (PSM) (Rosenbaum and Rubin, 1983) to estimate the distribution of rates of consumption. This is, to our knowledge, a novel application of PSM. Little (1986) applies propensity score weighting to missing data problems in sample surveys, but does not consider purchase infrequency.² The latter is not a missing data problem, since it arises irrespective of whether purchase diaries are fully completed. First we derive the theoretical case for using PSM for the imputation of consumption rates (section 2). We then estimate household motor fuel consumption rates using data from the UK's National Travel Survey (NTS) and evaluate the imputation statistically (section 3). We extend the analysis to study emissions reduction policies for household motor fuels in section 4.

2. Theory: using PSM to estimate consumption rates under infrequency of purchase

Let Z denote a binary event; $Z=1$ if a household purchased fuel and $Z=0$ otherwise. Let r denote a potential survey outcome, the quantity of fuel purchased conditional on $Z=1$. If $Z=1$, r is recorded, as the fuel purchase value in the dataset, otherwise 0 is recorded and r is unknown. We first estimate the missing values of r . We then use r in conjunction with propensity scores to estimate consumption rates, c .

A propensity score, $ps(\mathbf{X})$, is the conditional probability that Z occurs, given a vector of observed characteristics \mathbf{X} of a unit of observation i . Rosenbaum and Rubin (1983) show that the ps is a 'balancing score', meaning that the distribution of \mathbf{X} will tend to be the same for random samples of units with the same value of $ps(\mathbf{X})$, whether $Z=1$ or $Z=0$. That is,

$$\mathbf{X} \perp Z \mid ps(\mathbf{X}) \tag{1}$$

Balancing in this sense is a large sample property of ps . The true ps is always unknown and can only be estimated, for example using a binary regression model. Rosenbaum and Rubin (1983) also show that ps can be used to correct for certain kinds of selection biases. The usual application is in the

² Economists have developed the 'Infrequent Purchase Model' (IPM) (Deaton and Irish, 1984; Blundell and Meghir, 1987). The IPM estimates simultaneously a logistic regression equation of the purchase decision and a linear regression equation for consumption rates. We are concerned with a more basic task, namely estimating the distribution of consumption rates.

estimation of effect sizes for observational studies, to control for self-selection into a treatment group. Here in contrast we account for self-selection into the category of purchasers during a diary window. The key conditions required to estimate the sample distribution of r are

$$r \perp Z \mid \mathbf{X} \tag{2}$$

and

$$0 < p(Z=1 \mid \mathbf{X}) < 1 \text{ for all } \mathbf{X}.$$

(2) is known as ‘strong ignorability’. The first part means there are no unobserved confounders, that is, unrecorded variables that affect both the probability of purchase and the quantity r . The second part means that there is no \mathbf{X} such that Z is perfectly predictable. Given (2) it also follows that

$$r \perp Z \mid ps(\mathbf{X})$$

and

$$(3)$$

$$0 < p(Z=1 \mid ps(\mathbf{X})) < 1 \text{ for all } \mathbf{X}$$

From (1), estimated propensity scores, $\hat{ps}(\mathbf{X})$, of sufficient quality can always be used to balance samples on their observed characteristics. (3) implies additionally that each household i : $Z=0$, can be matched with a household j : $Z=1$ with approximately the same value of $\hat{ps}(\mathbf{X})$ to estimate an unobserved value of r :

$$\hat{r}_{i:Z=0} = r_{j:\hat{ps}_j(\mathbf{X}) \cong \hat{ps}_i(\mathbf{X}), Z=1} \tag{4}$$

where the relationship of proximity in estimated propensity scores, " \cong ", is operationalised by a matching algorithm.

The matched, purchasing households thus provide an estimate of the set of unobserved values of r . The quality of these estimates, given (2), will depend on both sample size and the quality of the estimated propensity scores. Here r represents the quantity purchased conditional on a purchase occurring in the diary window. We refer to r as the ‘quantity at the pump’ to distinguish it from the consumption rate, c . How long a household takes to consume a given quantity is inversely

proportional to its probability of purchasing. Each quantity is therefore multiplied by the corresponding estimated ps to yield an estimated rate of consumption, \hat{c} . That is, values given by

$$\hat{c}_i = \begin{cases} \hat{ps}_i(\mathbf{X}) \cdot r_i & \text{if } Z_i=1 \\ \hat{ps}_i(\mathbf{X}) \cdot \hat{r}_i & \text{if } Z_i=0 \end{cases} \quad (5)$$

constitute the estimated distribution of consumption rates. Although \hat{c} and \hat{r} are subscripted it is important to realise that a given imputed value is not an estimate *for that household*, since each value of the scalar $ps(\mathbf{X})$ is associated with a distribution of realisations of the vector \mathbf{X} , not a specific configuration. At each propensity score, that is, there is still heterogeneity, but it is unrelated to Z . PSM therefore results in group-level matching: a set of households is identified with a covariate structure which is expected to be identical to that of the $Z=0$ households.

Our application of PSM to infer the distributions of \hat{r} and \hat{c} is distinct from use of PSM for causal inference in observational studies. Firstly, in the latter context inferences from PSM generally only concern a mean, usually the mean effect of some intervention, rather than individual effects. This is because each effect is the difference between two potential outcomes, and one of these is unknown for each unit. Individual-level matching would be required to estimate individual effects and quantiles of the distribution. In the present setting, only one potential outcome is of interest, and it is unknown only for a subset of units. Secondly, in the causal inference problem the propensity score is only used to match units, whereas here it is used both for matching and to discount values for purchase infrequency. This implies a stronger condition for $\hat{ps}(\mathbf{X})$ to satisfy, since (2) can be satisfied even if there is omitted variable bias in the estimation of the probability of purchase.³ We therefore make explicit a distinct assumption:

$$\hat{ps}(\mathbf{X}) \text{ is an unbiased estimator of } p(Z=1) \quad (6)$$

(6) implies (2) and (3), since if (6) holds there is (for example) no omitted variable bias, so there are no unobserved determinants of Z correlated with \mathbf{X} , and therefore no such determinants of r and Z .

³ Suppose for example that all households purchased the same quantity (r) whenever they bought a particular good. Then (2) would be always be satisfied regardless of the quality of the propensity score model.

3. Estimating the distribution of fuel consumption using PSM and the UK National Travel Survey

3.1 Extent of infrequency of purchase

We consider data from the UK National Travel Survey (NTS), pooling data for years 2002-2008 to achieve a large sample size.⁴ The NTS is ideally suited for study of infrequent purchase for the following reasons. Firstly, given its diary window of one week many households do not purchase fuel. Secondly, it also records annual mileage for each vehicle in the survey interview, which provides a crude proxy for fuel consumption. Finally, the data concerned are policy-relevant, particularly for environmental and energy policy, so practical consequences of the data problem are salient.

Concerning the extent of infrequent purchase, the sample comprises a total of 57,069 fully-cooperating households. Of these, 42,712 have vehicles, either cars, vans or motorbikes, but 17,485 (41%) did not buy fuel during the diary week. Only 70 vehicle-owning households actually report zero annual mileage. So only around 0.2% of motoring households in the sample should have no fuel consumption and almost all the recorded zeroes result from infrequency of purchase. Histograms of the diary data and mileage data are shown in Figure 1 below. The diary data show a spike at zero and an extended tail to the right of the mean.

We have no reason to believe the mean purchase is a biased estimator of the mean consumption rate (n_2). Given the mileage data, however, the distribution of fuel consumption rates cannot resemble that in the left histogram of Figure 1. We anticipate a strong, direct relationship between the true distribution of mileage and the true distribution of fuel consumption rates. For, given the fuel efficiency of a vehicle, there is a determinate quantity of fuel required for a given journey. Consumption rates should therefore exhibit a distribution resembling that in the right histogram. The mileage data are not unproblematic,

⁴ The (unbanded) data are available on request through the Department for Transport.

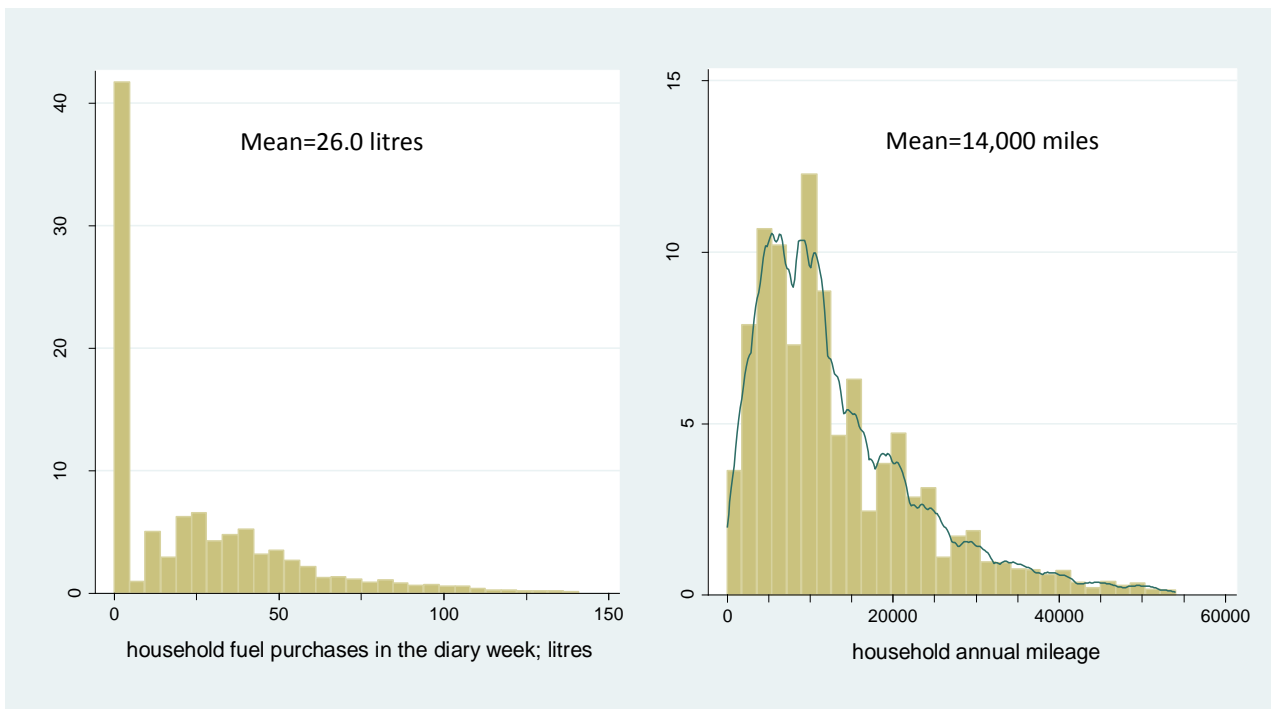


Figure 1. Sample distributions of household fuel purchases and mileage, NTS 2002-2008

Notes:

1. The NTS reports mileage separately for each vehicle. The figure is obtained by aggregating over vehicles.
2. Mileage is overlaid with kernel density estimates (Epanechnikov kernel).
3. Censored at the 99th percentile.

however, as the distribution has modes at multiples of 5000 miles, arising from over-reporting of salient numbers.

One could estimate fuel consumption directly from the mileage data, but there are serious disadvantages to doing so. Firstly, the NTS contains only discrete information relevant to the fuel efficiency of vehicles.⁵ For any given vehicle annual mileage there will in reality be a continuous distribution of fuel consumption rates. Secondly, the salient number bias would produce a multi-modal distribution of c . Our strategy instead is to use the mileage variable as one resource for matching-based estimation amongst other covariates.

Table 1 below gives more information about the occurrence of purchases in the sample. Whilst vehicle ownership is less common amongst less affluent households, the likelihood of non-purchase given ownership is higher. This implies that the divergence between the sample distribution of fuel purchases and that of the latent variable c is greater amongst less affluent

⁵ The relevant variables are a binary indicator of engine size (>1500 cc), a fivefold categorisation of vehicle type, and fuel type (diesel versus petrol) for each vehicle.

motoring households. However, the problem is pronounced everywhere. Amongst the top income quintile, for example, 1/3 of motoring households have no recorded purchase and purchases exceed weekly consumption rates by a mean factor of ~ 1.5 ($=1/(1-0.335)$).

Quintile of equivalised income	% of motoring households	% of motoring households with no fuel purchase
1	35.9	56.3
2	69.4	49.1
3	83.7	42.1
4	93.2	35.4
5	94.3	33.5

Table 1. Extent of infrequency of purchase by income quintile

Notes:

1. motoring households are defined as those owning at least one motor vehicle (car, van or motorcycle).
2. equivalised income is calculated using the 'square root scale' (OECD 2013) throughout.

3.2 Using PSM to recover 'quantities at the pump'

The PSM is conducted using the vehicle-owning households only. Matching with replacement was applied, with a caliper of 0.01, using the `psmatch2` routine in STATA (Leuven and Sianesi, 2003). In this approach, \hat{p}_s is the fitted value of a probit regression model. Each household which did not buy fuel is matched using \hat{p}_s to one that did, but the same match can be used more than once. This procedure is heterogeneity-preserving, which is appropriate here since we are attempting to recover an entire distribution.

Two probit models were developed. The results are shown in a coefficient plot (Jann, 2013) in Figure 2 below. Model 1 makes full use of relevant covariate information in the NTS excluding the annual mileage variable. Although the NTS has not been collected to estimate fuel purchase propensity, it provides a rich set of relevant variables. We exclude mileage to see how the PSM-based imputation fares in the absence of a proxy for the imputed variable, since this will be the usual

research situation. Square terms for age and numbers of adults are included, plus an interaction term for working households with




Figure 2. Coefficient plot of probit regression models of fuel purchase in the diary week

Notes

1. Probit coefficients are shown as diamonds with lines representing confidence intervals.
2. Continuous regressors are standardised to have mean zero and unit variance.
3. Household-level variables are derived from individual- and vehicle-level data (authors' calculations).
4. A dummy for each top-coded variable is included in the estimations but not shown.
5. For model 1 $r\text{-sq}=0.055$ (McFadden's pseudo $r\text{-sq}$) and $\text{Log-L} = -27250.28$. Model 2 has 2 additional parameters; $r\text{-sq}=0.080$ and $\text{Log-L} = -26534.82$. The LLR statistic is therefore $1430.92 \sim \chi^2(2)$; $p < 0.001$.

children, since these were found to improve goodness of fit and matching quality. Model 2 simply adds the mileage variable as a regressor.

Figure 3 below shows the distributions of \hat{ps} in models 1 and 2, respectively, using kernel density plots. Common support, the second part of condition (3), is approximately satisfied. There are no unmatched households under model 1, but under model 2, 3 households are unmatched because of the 1% \hat{ps} caliper we apply, and dropped. We regard this proportion as negligible.

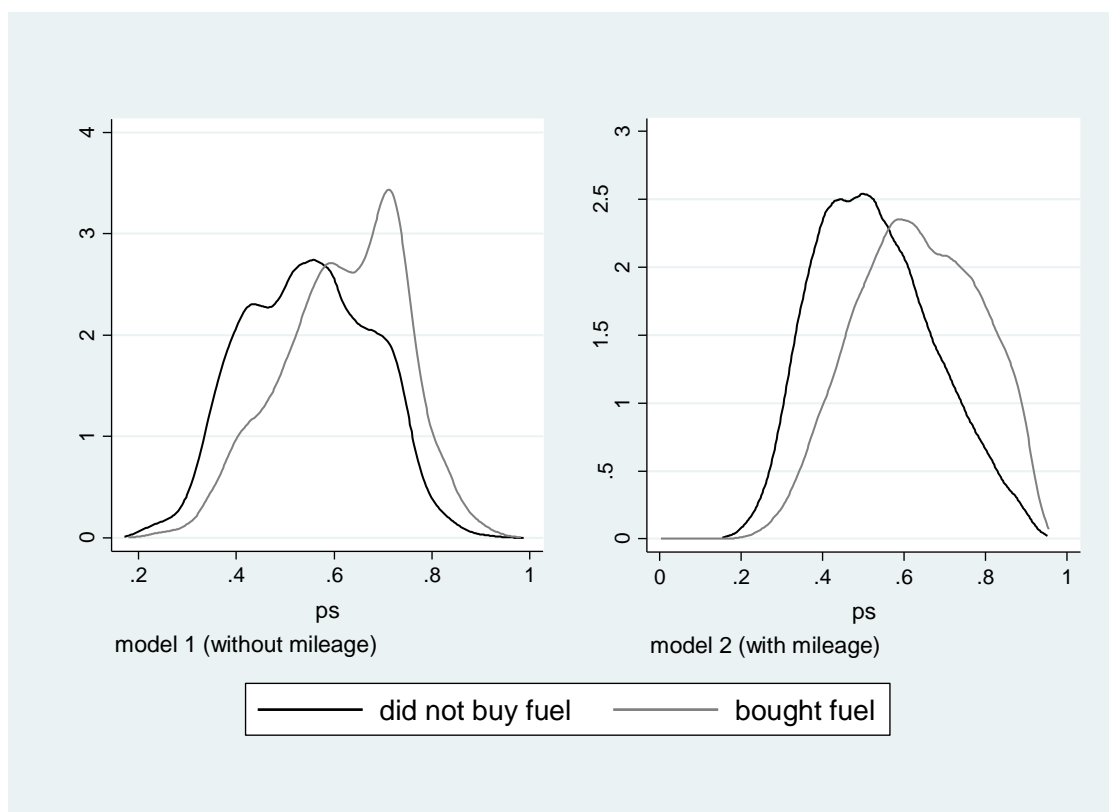


Figure 3. Estimated propensity scores: kernel density estimates

Note: Epanechnikov kernel

The pronounced multi-modal distribution of model 1 propensity scores (Figure 3) is attributable to particular constellations of covariates with high-valued regression coefficients: two adult, two car, rural households with children, for example. Within such groups the distributions are approximately unimodal.

Model 1 shows results which are generally in line with expectations, with positive coefficients for the number of driving licence holders, adults and children, the number of vehicles, distance from a train station and rural location, for example. Negative coefficients for diesel and motorcycles presumably reflect fuel efficiency. Model 1 performs poorly though, in terms of balance between the matched groups on the annual mileage variable. A visualisation of covariate balance is provided in Figure 4 below. Standardised percentage bias (Rosenbaum and Rubin, 1985) is shown before and after matching for each coefficient, in order of pre-matching bias. Covariates that are included in the regression have very low

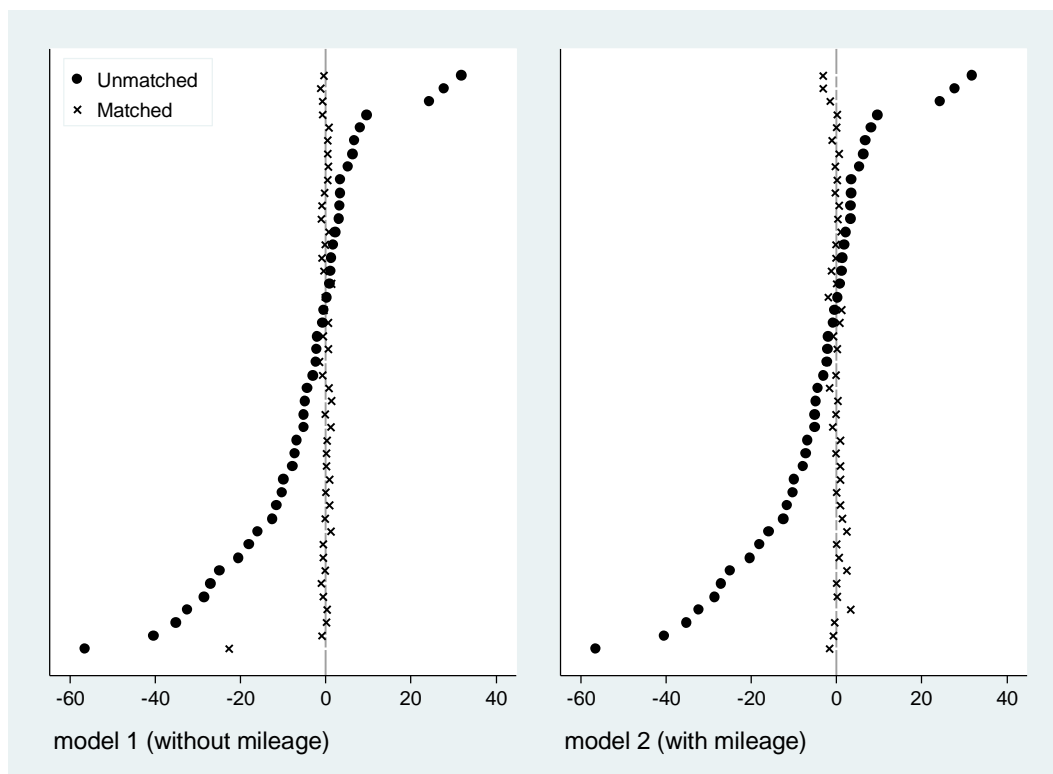


Figure 4. Standardised percentage bias between covariates in Z=0 and Z=1 households

Notes

1. Covariates are those shown in Figure 2, less square and interaction terms, plus topcode dummies.
2. The mean, median and maximum absolute standardised % bias are 1.1, 0.6 and 22.7 for model 1 and 1.0, 0.7 and 3.3 for model 2.

standardised biases (less than 2%).⁶ However, mileage, the coefficient at the bottom, shows the highest bias, exceeding 20%. Using \hat{ps} from model 1, therefore, we obtain matched groups with

⁶ Austin (2011) reports that standardised bias of less than 10% are regarded as low in applied work.

significantly different mean mileage and, therefore, systematically different actual ps . This violates our requirement (6).

In model 2, the coefficient on mileage dominates the regression. Given the physical relationship between mileage and fuel consumption this is unsurprising. Since many of the independent variables are determinants of mileage, some coefficients change sign or become insignificant. Figure 4 indicates that the two groups are now well-matched on mileage, with only a slight worsening of the bias metric on the other independent variables.

Comparing (nested) models 1 and 2, conventional model selection criteria also favour model 2 (note 5 to Figure 2), which is therefore our preferred means of estimating ps .

Having constructed the matched groups using PSM (with model 2), we take values for \hat{r} from the matched set of $Z=1$ households as stipulated in equation (4). Thus, for households who were observed buying fuel, we have observations of r and for $Z=0$ households we have PSM estimates of quantities they would have bought, had they made a purchase, \hat{r} . Values of r and \hat{r} are shown in Figure 5 below.




Figure 5. Quantities at the pump (litres) derived from PSM using model 2

Figure 5 shows greater frequency of lower quantities amongst non-purchasing households. This is consistent with the difference in propensity scores between the two groups (Figure 4), and an association between Z and r prior to controlling for X . Also noticeable is the pronounced multimodality of the distribution, with modes at multiples of 10 litres. Presumably this reflects a combination of over-reporting, and actually purchasing, salient numbers. Modes at 12-13 and 24-26 litres may be explained as follows. From 2002-2005, the price of petrol was roughly £0.80p per litre (AA, 2002-2009). Thus, each £10 spent on petrol would result in a purchase of around 12.5 litres for half the period under consideration.

Figure 5 also illustrates the heterogeneity-preserving quality of the matching-based imputation procedure. The same pattern of modes at salient numbers is evident for both observed and imputed purchases.

3.3 Estimated fuel consumption rates

Having derived quantities at the pump, the next step is to multiply each quantity by its associated propensity score to obtain estimated consumption rates, \hat{c} , as specified in equation (5). The resulting estimates are shown in Figure 6 and summarised in Table 2 below, alongside estimates using \hat{p} s from model 1, the diary fuel purchase and annual mileage variables.

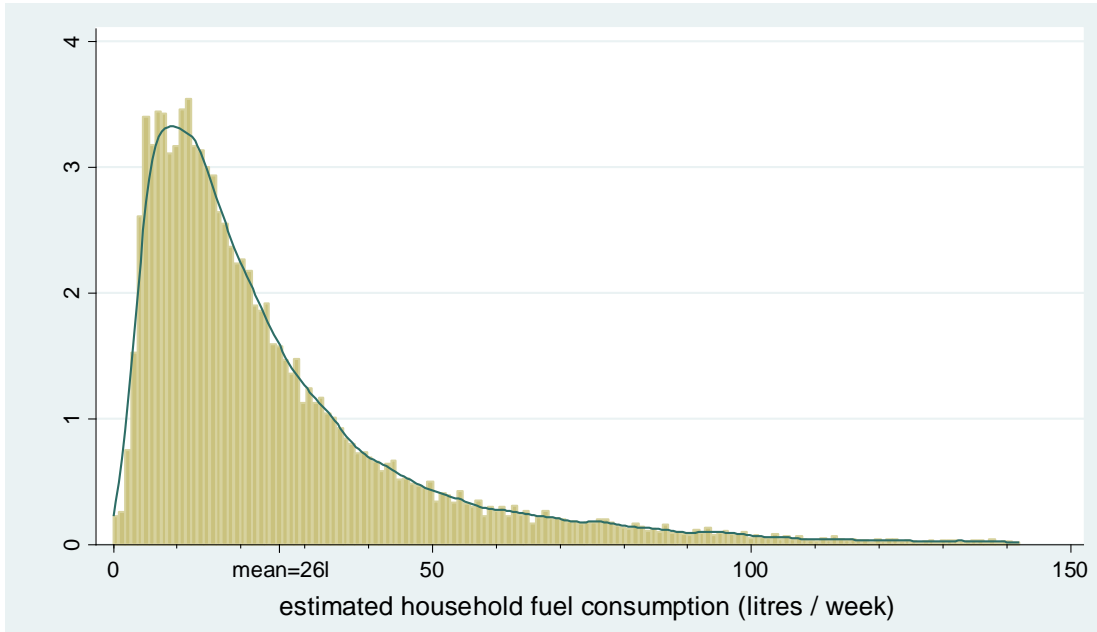


Figure 6. Preferred estimates of fuel consumption rates, derived from model 2

Notes

1. Kernel density estimates (Epanechnikov kernel) overlaid
2. Excludes top percentile

Standard errors for fuel consumption in Table 2 are calculated by bootstrapping, incorporating variation associated with the \hat{p}_s estimation and matching. Regular bootstrapping in this context fails to reproduce the distribution of times a unit is used as a match, f_i (Abadie and Imbens 2008). We avoid this problem by adding a small random error, e , to \hat{p}_s after drawing each bootstrap sample but before conducting the \hat{p}_s matching. This resolves the problem in theory, given values of e small enough for equation (4) still to hold but large enough to perturb the match selected. The distribution used was $e \sim N(0, 1/30625)$.⁷

We offer the following observations on the quality of the preferred estimates (Table 2, column 4). \bar{c} approximately equals the mean fuel purchase (26.06 litres versus 26.03 litres

⁷ We selected parameters for e which approximately reproduce the distribution of f_i without detriment to the standardised bias metric of matching quality, by trial and error (Table A1, Appendix). We also tested our bootstrap procedure using Monte Carlo simulation (Table A2, Appendix). The bootstrap standard errors for the mean and quantiles of the distribution approximate standard deviations of the corresponding variables derived using simulated samples, but those for standard deviation and skewness do not, a problem which seems attributable at least in part to skewness of c (notes to Table A2, Appendix). We therefore include standard errors only for the mean and percentiles in Table 2.

respectively), as required. At the same level of granularity, the multimodality of Figure 5 is absent from Figure 6, which is reassuring since it is unlikely that c is affected by salient number biases. The distribution also appears plausible judging the mileage proxy. Let Q_1 , Q_2 and Q_3 denote the 25th, 50th and 75th percentiles of a distribution respectively. The proportional relationships Q_1/Q_2 and Q_1/Q_3 are identical for mileage and \hat{c} to one decimal place. For a more detailed comparison we present quantile-quantile plots in Figure 7, normalising by dividing each value by the maximum of the variable.

statistic		reported fuel purchase in diary week (litres)	reported annual mileage (thousand miles)	weekly fuel consumption from PS model 1	weekly fuel consumption from PS model 2 (preferred)
	1	0	0.7	3.1 (0.11)	2.6 (0.10)
	5	0	2	5.1 (0.08)	4.6 (0.08)
	10	0	3	6.7 (0.09)	6.1 (0.09)
	25	0	6	11.2 (0.12)	10.4 (0.12)
percentile	50	18	10	19.1 (0.15)	18.3 (0.15)
	75	40	18	32.7 (0.26)	32.5 (0.25)
	90	67	28	54.5 (0.55)	55.3 (0.51)
	95	90	35	71.2 (0.75)	75.6 (0.74)
	99	142	54	115.1 (1.77)	124.5 (2.13)
mean		26.0	13.7	26.1(0.18)	26.1 (0.16)
std		33.1	11.3	23.4	25.2
skewness		2.4	2.1	2.9	3.0
N		42600	42707	42598	42595

Table 2. Estimates from PSM based imputation compared to recorded fuel purchase and annual mileage

Notes

1. Bootstrap standard errors in parentheses, with 1000 repetitions. See main text and note 8 for discussion.
2. N varies across columns because of missing NTS data. In addition, 3 observations cannot be matched using model 2 under the 1% *ps* caliper restriction applied.




Figure 7. Quantile-quantile plots of estimated fuel consumption rates against recorded annual mileage.

Note

Values have been divided by the maximum of the range.

The left plot of Figure 7 shows that the quantiles of fuel consumption are located somewhat lower in their range, than are quantiles of mileage.⁸ Since for both variables the 99th percentile is less than 0.35 of the maximum value, the right plot is drawn for percentiles 1-99 only. This confirms that values are somewhat more concentrated at lower areas of the range for fuel consumption, consistent with the difference in skewness shown in Table 2. This may be associated with features of the distribution of vehicle fuel efficiency. However, in both cases the plots do not deviate dramatically from the 45-degree line and the larger deviation concerns the top 1% of observations.

⁸ For example, Q3 fuel consumption is .07 of the maximum value (32.5/469). Q3 mileage is .12 of the maximum value (18,000/153,000). So one point in the above plot is (.07, .12).

3.4 Estimated fuel consumption rates without mileage in the propensity score model

The corresponding quantile-quantile plots using estimates from model 1 are very similar to Figure 7. They show a greater deviation from the 45 degree line for the full range plot, and less deviation from it for the first 99 percentiles (Figure A1, Appendix). From Table 2, although the percentiles obtained under the two models are generally significantly different, this is attributable to the large sample size. The absolute differences in \hat{c} are fractions of a litre per week excepting at the upper tail and the mean is approximately the same. Thus, even without the mileage proxy included in the ps estimation the estimates are plausible.

A detailed exploration of the robustness of the method is beyond the scope of this paper, but we offer the following observations. PSM using model 1 matches households with systematically different actual fuel use, since they have systematically different mileage (Figure 4, left). Ignorability, (3), may approximately hold without including mileage in \mathbf{X} , however, since it is a condition on r , not c . Consistently with this, the distributions shown in Figure 5 are very similar if we use model 1 estimates (Figure A2, Appendix; the estimates differ by 1l at the median and 1.2l at the mean, about 3% in each case).

It therefore seems possible that under model 1, draws from the same distributions of quantities occur when the algorithm pairs $Z=1$ and $Z=0$ households. But they are then multiplied by the same value of \hat{ps} when different values of ps actually obtain. This is consistent with the higher standard deviation of \hat{c} under model 2 (Table 2), and the larger variance in \hat{ps} (Figure 3). It seems from the plausibility of the model 1 distribution, however, that the resulting errors are counterbalanced to a significant degree. This may not be surprising, given that with an appropriately designed and implemented sample survey we have an unbiased estimate of \overline{ps} ($=\bar{Z}$) prior to any modelling of the purchase decision. The case studied therefore seems encouraging from the perspective of the 'applied researcher', who realistically lacks a proxy for the target variable. For it seems that omitted variable bias in the ps model may sometimes have little effect on the bulk of the estimated distribution. However, the two sets of estimates are more divergent at the highest

quantiles of estimated consumption. This may be because model 1 underestimates the occurrence of households for whom $ps > 0.75$ (Figure 3).

4. Application to emissions reduction policy

4.1 Estimation of UK household CO₂ emissions from motor vehicles

Given increasing greenhouse gas concentrations in the atmosphere, it is interesting to consider the relevance of our results to discussions of household CO₂ emissions, particularly since infrequent purchase has constrained their analysis.⁹ We calculated CO₂ emissions of each vehicle using the fuel purchase diary and DECC / DEFRA emissions factors (DEFRA, DECC 2011), using separate figures for petrol and diesel. These figures were then aggregated to yield motoring emissions for each household. The resulting estimates suffer from essentially the same infrequency of purchase problem outlined above, and are treated in the same way. That is, we substitute the emissions quantity for each $Z=0$ household with the value obtained for its ps -matched observation, using model 2 estimates, and then multiply each emissions quantity by its estimated ps . The resulting estimates, \hat{a} , are strongly isomorphic to \hat{c} , since emissions are simply a multiple of the amount of each fuel purchased representing its carbon content. Mean (median) annual motoring emissions over the period are calculated to be 2.4 (1.5)t CO₂ per household, or 3.2 (2.2)t CO₂ per motorist household.

Of particular interest is the estimated concentration of emissions, a notable study having reported that they are disproportionately accounted for by a relatively small group of high-emissions households (Brand and Boardman, 2008). We summarise estimated shares of vehicle CO₂ emissions by (emissions) decile in Table 3 below.

⁹ On UK households' greenhouse emissions, see for example Gough *et al.* (2011) and Büchs and Schnepf (2013).

Emissions Decile	% share (SE)	Cumulative % share
1	1.7 (.03)	1.7
2	2.9 (.04)	4.5
3	4.1 (.05)	8.4
4	5.2 (.05)	13.4
5	6.4 (.06)	19.6
6	7.9 (.07)	27.3
7	9.8 (.08)	37.0
8	12.5 (.11)	49.5
9	17.3 (.16)	66.9
10	32.4 (.34)	100

Table 3. Decile shares of UK household CO₂ emissions from motor fuels

Note. Bootstrap standard errors in parentheses, with 1000 repetitions. See main text and note 8 for discussion.

This breakdown confirms the concentration of vehicle emissions, with an estimated 1/2 of motor fuel CO₂ accounted for by the top quintile, and 1/3 by the top decile alone. The advantage of our estimates is that they are based on a national representative survey. Brand and Boardman used a local sample survey conducted in Oxfordshire coupled with an online survey, so the estimates have an ambiguous geographic and statistical status. The authors also report a ratio between the top and bottom quintiles of 15:1. Our estimate for the UK is lower, but still remarkable, at 10.9:1 with 95% c.i. $(10.5 \leq x \leq 11.2):1$ ($\pm 1.96 \times$ bootstrap standard error).

As Brand and Boardman (2008) suggest, the policy implication of the high concentration of emissions is that reducing those of a relatively small proportion of (generally richer) high emitting households would be highly effective in terms of tackling overall emissions. In absolute terms the policies usually discussed, namely carbon taxes and carbon rationing, would both affect higher emitting households more, but operate regardless of income *per se*. How such policies would affect different income groups has therefore attracted much attention (Büchs *et al.*, 2011). The literature

has been unable to estimate the spread of policy impacts *within* different income bands, however, since the available national surveys are all affected by infrequent purchase. So although mean effects have been estimated by income group, it is not known how representative these are. That they may be heavily influenced by relatively extreme values is suggested by the high skewness of \hat{c} in Table 2. For further insight we use our estimates of \hat{a} and covariate information in a simple simulation of emissions reduction policy.

4.2 Motor fuel emissions reduction policies

4.2.1. Static microsimulation

Ideally one would conduct a sophisticated policy simulation incorporating behavioural responses and a model of the economy (Barker, 1998). Examples include REMI (2014) for emissions reduction policies for the USA, and Comhar (2008) for transport fuel policies in the Republic of Ireland for example. But that would constitute a complex study in its own right, and introduce many additional sources of uncertainty. Instead we use the simplest approach, static microsimulation, which assumes behaviour is unchanged, to illustrate directly the value of our method. In essence, this calculates ‘who stands to lose how much’ under a policy. Such analysis is frequent in the media and offers a starting point for policy evaluation. It offers insights into policy resistance / acceptance, is an important step towards analysis of regressivity, and provides insight into probable early effects, since behaviour and the economy take time to change.

For recent discussion of emissions reductions policies, including implementation issues, see Sorrell (2010). Two such policies are considered here. The first is a carbon tax or tradeable ration / cap. Taxes and caps would have very important differences in practice,¹⁰ but are analytically equivalent within our framework. Thus, we can assume that a tax is levied at a certain rate without behavioural response, or that fuel use is capped at current levels and price responds by an assumed amount because of scarcity at the margin as consumers have to buy permits. We will assume

¹⁰ Basically taxes have uncertain impact on emissions, whereas a cap has uncertain effects on prices. See Sorrell (2010) for a recent proposal including discussion of implementation issues.

£100/tCO₂ as the tax rate (or price increase). The second policy is the same tax or ration implemented in revenue-neutral form. That is, the carbon revenue is allocated to the households on an equal per-capita basis, with each adult aged 16 or over allocated an equal share. Tax and ration / cap variants of this policy are known as ‘tax and dividend’ (Hansen, 2009) or ‘cap and share’ (Comhar, 2008) respectively. A household’s net payment, v , is defined as its payment for carbon content of its fuel, t , minus its income from per-adult revenue shares.¹¹

Since our estimates are derived using matched groups, we cannot identify a Z=1 household a given Z=0 household is matched to. We need household-level information though, to calculate outcomes by income decile and / or as proportions of income. We address this problem exploiting property (1) of ps : the structure of observed covariates in the matched controls, for a large sample and true ps , is identical in expectation to that in the original group. Income is an observed covariate, included as a regressor to calculate \widehat{ps} . We therefore substitute the Z=0 households for the matched controls, dropping the former from the dataset. In the next 2 subsections we report estimates of (net) payments under the two policies, expressed in absolute terms and as proportions of income:

¹¹ The analysis concerns only revenue collected on household vehicles, since these constitute the NTS dataset. Thus, business transport fuel is assumed to fall outside the scope of the policy and we are not able here to include indirect effects that would result from their inclusion.

$$\hat{t}_i = 100 \hat{d}_i = \text{tax} \quad [\text{£}] \quad \text{and} \quad \frac{\hat{t}_i}{y_i},$$

and

$$\hat{v}_i = \hat{t}_i - pm_i = \text{rebated tax (net payment)} \quad [\text{£}] \quad \text{and} \quad \frac{\hat{v}_i}{y_i},$$

where

$$p = \frac{100 \sum_{i=1}^n \hat{d}_i}{\sum_{i=1}^n m_i} \quad [\text{£}] = \text{the value of the per capita revenue share,}$$

m_i = the number of persons aged ≥ 16 in a household, and

y_i = equivalised household income.

4.2.2 Effects of a carbon tax or ration of motor fuels on UK households

\hat{t}_i and $\frac{\hat{t}_i}{y_i}$ are shown using distribution plots in Figure 8 below. The plots show the 10th and 90th percentiles, Q1, Q2, Q3 and means of the estimate over quintiles of equivalised household income. The two measures of central tendency are shown connected to show the gradient across income quintiles. Although we stipulate a £100/tCO₂ carbon price, since \hat{t} is simply a multiple of \hat{d} , estimated effects at other prices can be directly inferred. For example, at £200/tCO₂ each figure on the y-axes would be doubled.

A clear feature of \hat{t} evident in the left-side plots is its consistent positive skew. For motorists, the mean payment exceeds Q2 by 34-40%, and in the first two quintiles is closer to Q3 than Q2. Thus, reporting on mean effects considerably amplifies the impact on a ‘typical’ household compared to the more representative and robust Q2 and seems particularly misleading at lower incomes. Another interesting feature of the distributions is the relatively low gradient between the 1st and 2nd income quintiles of motorists (2nd plot from the left). This may be an indication that car use is mainly for essential journeys rather than leisure at lower incomes. The higher gradient between these quintiles in the leftmost plot reflects the increase in vehicle ownership with income.

The right hand plots show $\frac{\hat{t}_i}{y_i}$. The tax appears regressive amongst motorists (rightmost plot), with the gradient in incomes exceeding that in absolute payments. However, the mean proportional tax is closer to Q3 than Q2 for quintiles 1-3 indicating heavy influence by relatively extreme values. For the 1st quintile, for example, it is approximately double the Q2 value of ~3%.




Figure 8. Estimated monetary effects of a carbon tax or ration at £100/tCO₂

Note: ‘diary sample’ weights applied, recalculated for the pooled sample.

The picture is more complicated for the population as a whole (2nd plot from the right), due to low rates of vehicle ownership in the 1st and 2nd quintiles. At the mean, the policy is estimated to be regressive, though $\frac{\hat{t}_i}{y_i}$ does not decline monotonically across quintiles.¹² We estimate that

¹² Our results at the mean contrast with an earlier claim in the policy literature, that taxes on motor fuels are progressive overall, and only regressive amongst motorists (Blow and Crawford, 1997; Dresner and Ekins, 2004). This difference is likely to be attributable to increasing car ownership over time. According to the NTS, 44% of households in the lowest income quintile owned or rented a car over 2002-2008, up from 34% in 1995/1997 (DFT, 2012 and own calculations).

evaluated at the median, the tax is progressive across quintiles 1-3 but slightly regressive across 3-5. Again, the mean proportional tax is closer to Q3 than Q2 for quintiles 1-3.

4.2.3 Effects of 'cap and share' or 'tax and dividend' for motor fuels on UK households

\hat{v}_i and $\frac{\hat{v}_i}{y_i}$ are represented in Figure 9 below for a tax rate of £100/tCO₂ rebated to the population.

The value of the per capita payment at this tax rate is estimated at £127. Again, effects of different CO₂ prices can be directly inferred by rescaling the axes. In these plots, in addition to the gradient, it is interesting to consider the predicted proportions of the quintiles or population that stand to win ($\hat{v} < 0$) and lose ($\hat{v} > 0$) financially.

Consider first the left-side plots, showing \hat{v}_i . Consistently with Figure 8, there is a low gradient between quintiles 1 and 2 in terms of absolute payments amongst motorists, reflecting broadly similar patterns of car use. The gradient between these quintiles in the leftmost plot (compared to the corresponding figure in Figure 8) is lower. This reflects fewer single-adult households in quintile 2 than in quintile 1 (39% versus 58%), and consequently more dividend payments, compensating for increased car ownership. There is again positive skew but less extreme than in Figure 8, with means slightly closer to Q2 than to Q3. Considering winners and losers, the mean household gains in quintiles 1-3 (all households) or 1-2 (motorist households), whereas the median household gains in quintiles 1-4 or 1-3

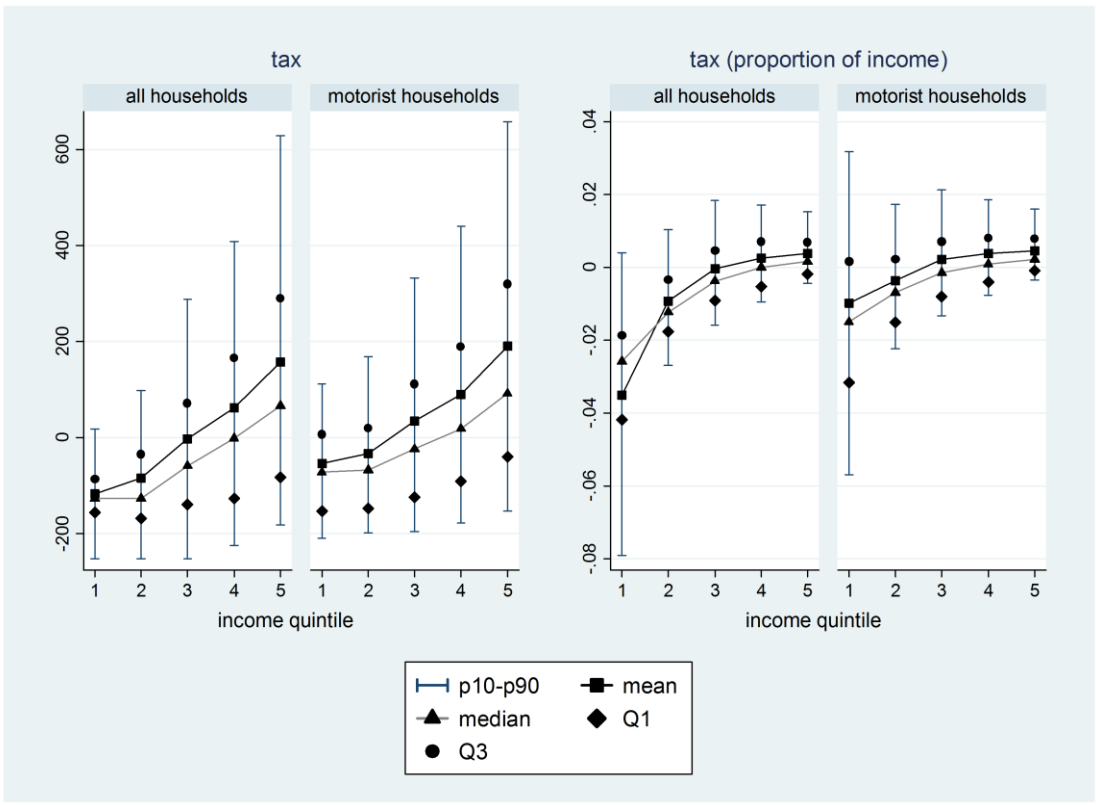


Figure 9. Estimated monetary effects of a 'cap and share' or 'tax and dividend' policy at £100/tCO₂
 Note: 'diary sample' weights applied, recalculated for the pooled sample.

respectively. Overall, only the richest quintile are estimated more likely to lose than gain, and most motorist households are actually estimated to gain. Analysis at the mean conceals these features, which are highly politically salient.

Now consider the right-side plots, showing $\frac{\partial \ln Y_i}{\partial \tau_i}$. In contrast to Figure 8, these show strongly progressive outcomes, with relatively large percentage gains at lower incomes, paid for by relatively small transfers from higher income households. In quintile 1 (all households), there is negative skew, which occurs because a relatively small group of households do extremely well proportionally under the policy: namely those without vehicles with very low cash incomes and large numbers of over-16s. The mean here is closer to Q1 than Q2. The other distributions exhibit positive skew. For example, the mean household in quintile 1 stands to gain an estimated 1% of income and the median household an estimated 1.5%, a proportional gain 50% higher.

Lacking estimates of quantiles because of infrequency of purchase may therefore have appreciable impact on policy analysis. Previous studies of emissions reduction policies for motor fuels had to rely on mean consumption rates, including means for different income bands. But a carbon tax or cap appears considerably more regressive, and its revenue-neutral counterpart less progressive, when evaluated at the mean rather than at the median. In short, the policy seems better with the quantile information. The picture is not uniform however, since not all the estimated distributions exhibited positive skew. We judge that the overall pattern could not be predicted simply by inspection of the underlying variables.

The p10-p90 ranges for the lowest income households in the rightmost plots of Figures 8 and 9 suggest extreme heterogeneity, which could be politically problematic since extreme cases often receive prominent media attention. These estimates are probably affected by further data limitations, however, since low income households may rely heavily on the benefits system, which is not accounted for in the NTS. Additional data collection would presumably be necessary to better evaluate outcomes at the lowest incomes.

Finally, we note that conducting the same policy simulation using propensity scores from model 1 to estimate \hat{a} produces almost identical results across the bulk of the distribution. The graphs obtained corresponding to Figures 8 and 9 are visually distinguishable only at the 90th percentile for quintiles 4 and 5 in the left side plots (Appendix, Figures A12 and A13). This is consistent with our earlier observation that estimates from the two models of fuel consumption rates differ substantively only in the right tail of the distribution.

5. Conclusions

A simple method was presented whereby propensity scores can be used to adjust a variable affected by a short observation window in sample surveys, a longstanding problem precluding distributional analysis. First, match each Z=0 unit to a Z=1 unit on the \widehat{ps} to obtain \hat{r} for the former. Second, multiply each value of r (if Z=1) and \hat{r} (if Z=0) by \widehat{ps} to obtain estimates of the latent variable c of

interest. The problem and method were illustrated using the UK National Travel Survey, which contains a proxy (annual mileage) for the affected variable (fuel purchase). The resulting estimates of fuel consumption rates are plausible judging by the distribution of household annual mileage calculated from the same survey. Estimates obtained without recourse to the mileage proxy are also plausible, differing substantially from our preferred estimates only in the upper tail of the distribution. This is encouraging, since a proxy for the target variable will not normally be available.

The method was then applied to conduct a static microsimulation of two emissions reduction policies for motor fuels, supposing a carbon price of £100/tCO₂. Such exercises have previously had to rely on estimates of mean effects. We judge that estimating entire distributions of effects shows the policies in a more favourable light. The distributions appear to be highly skewed, influencing the mean appreciably, but not always in a consistent direction. This information is timely given the outcome of the recent COP 21 meeting, which agreed targets and aims for curtailment of global warming, but did not agree any emissions reduction policies to achieve these.

A simple carbon tax or ration / cap would be regressive amongst motorists, but appears less regressive evaluated at the median than at the mean. The same policy conducted in revenue-neutral form, for example by redistributing revenues on an equal per-capita basis, is estimated to benefit the majority of households in all but the top income quintile, and even the majority of motorist households overall. These important features of the policy are hidden under analysis at the mean. The gains would result from a relatively small estimated transfer from generally higher income households. This is because of the high concentration of estimated motoring emissions and their strong association with income.

Our estimation of who stands to lose and gain financially raises a key question for comparison of 'tax and dividend' and 'cap and share' variants of the revenue-neutral policy. Since the marginal propensity to consume varies inversely with income,¹³ it cannot be ruled out *a priori*

¹³ For example, using Italian data, Japelli and Pistaferri (2013) report that the poorest households sampled spend on average c.70% of additional income, whilst the richest spend only c.35%, and estimate that

that fuel consumption would increase under tax and dividend, contrary to the environmental goal, despite a higher fuel price. To address that issue requires going beyond static microsimulation. Either policy would plausibly increase consumption expenditure overall.

A limitation of the method presented is that one needs to know whether a unit records a zero value because of a short observation window or for some other reason. In the NTS one can distinguish between infrequency of purchase and non-consumption of motor fuels, because vehicle ownership is recorded. In other consumption surveys, including the widely-used Living Costs and Food Survey, this is not known for many items. So for a long time econometricians have endeavoured to distinguish between non-consumption of meat, tobacco and alcohol, for example, and infrequency of purchase. An implication of the present study is that inclusion of a question to identify non-consumption of important items has potentially large benefits at relatively small cost. Since this may allow researchers to apply simple matching methods to recover the distribution.

References

- AA (2002-2009). Fuel Price Reports. The Automobile Association.
- Abadie, A. and Imbens, G. (2008). On the failure of the bootstrap for matching estimators. *Econometrica*, 76, 1537-1557.
- Austin, P.C. (2011). An introduction to Propensity Score Methods for Reducing the Effects of Confounding in Observational Studies. *Multivariate Behavioural Research*, 46, 399-424.
- Barker, T.S. (1998). Use of Energy-Environment-Economy Models to Inform Greenhouse Gas Mitigation Policy. *Impact Assessment and Project Appraisal*, 16, 123-131.
- Blow, L. and Crawford, I. (1997). The distributional effects of taxes on private motoring. Institute for Fiscal Studies: London.

transferring 1% of national income from the richest to the poorest income decile would increase consumption expenditure by 0.33% .

Blundell, R. and Meghir, C. (1987). Bivariate Alternatives to the Tobit Model. *Journal of Econometrics*, 34, 179-200.

Brand, C. and Boardman, B. (2008). Taming of the few: the unequal emissions of greenhouse gas emissions from personal travel in the UK. *Energy Policy*, 36, 224-38.

Büchs, M., Bardsley, N. and Duwe, S. (2011). Who bears the brunt? Distributional effects of climate change mitigation policies. *Critical Social Policy*, 31, 285-307.

Büchs, M., and Schnepf, S.V. (2013). Who emits most? Associations between socio-economic factors and UK households' home energy, transport, indirect and total CO2 emissions. *Ecological Economics*, 90, 114-123.

Burton, A.C., Neilson, E., Moreira, D., Ladle, A., Steenweg, R., Fisher, J.T., Bayne, E. & Boutin, S. (2015). Wildlife camera trapping: A review and recommendations for linking surveys to ecological processes. *Journal of Applied Ecology*, 52, 675-685.

Cohmar Sustainable Development Council (2008). A study in personal carbon allocation: cap and share. Consultancy report prepared by AEA consulting and Cambridge Econometrics. Comhar: Dublin.

Deaton, A.S. and Irish, K. (1984). Statistical models for zero expenditures in household budgets. *Journal of Public Economics*, 23, 59-80.

DECC, DEFRA, (2011). 2011 Guidelines to Defra/DECC's GHG conversion factors for company reporting. Department for Environment, Food and Rural Affairs and Department for Energy and Climate Change: London

DFT (2012). Table NTS0707: Adult personal car availability by ethnic group: Great Britain. Department for Transport: London.

Dresner, S., and Ekins, P., (2004). The distributional impacts of economic instruments to limit greenhouse gas emissions from transport. PSI Research Discussion Paper 19, Policy Studies Institute: London.

Gough, I., Abdallah, S., Johnson, V., Ryan-Collins, J. and Smith, C. (2011). The distribution of total embodied greenhouse gas emissions by households in the UK, and some implications for social policy. CASE PaperNo.152, London: Centre for the Analysis of Social Exclusion, London School of Economics.

Hansen J.E. (2009). *Storms of My Grandchildren*. Bloomsbury: NY.

Jann, B. (2013). Plotting regression coefficients and other estimates in Stata. University of Bern Social Sciences Working Papers Nr. 1.

Jappelli, T. and Pistaferri, L. (2013). Fiscal policy and MPC heterogeneity. *American Economic Journal – Macroeconomics*, 6, 107-136.

Leuven. E. and Sianesi, B. (2003). PSMATCH2: Stata module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing. <http://ideas.repec.org/c/boc/bocode/s432001.html>

Little, R.J.A. (1986). Survey Nonresponse Adjustments for Estimates of Means. *International Statistical Review*, 54, 139-157.

OECD (2013). What are equivalence scales? (Online document available at <http://www.oecd.org/els/soc/OECD-Note-EquivalenceScales.pdf>, downloaded 02/01/2015)

REMI (2014). The economic, climate, fiscal, power, and demographic impact of a national fee-and-dividend carbon tax. Regional Economic Models Incorporated: Washington DC.

Rosenbaum, P.R. and Rubin, D.B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 41-55.

Rosenbaum, P.R. and Rubin, D.B. (1985). Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician* 39, 33-38.

Sorrell, S. (2010). An upstream alternative to personal carbon trading. *Climate Policy*, 10, 481-486.

Appendix




Figure A1. Quantile-Quantile plots of estimated fuel purchases against annual mileage; comparison of model 1 and model 2 estimates.

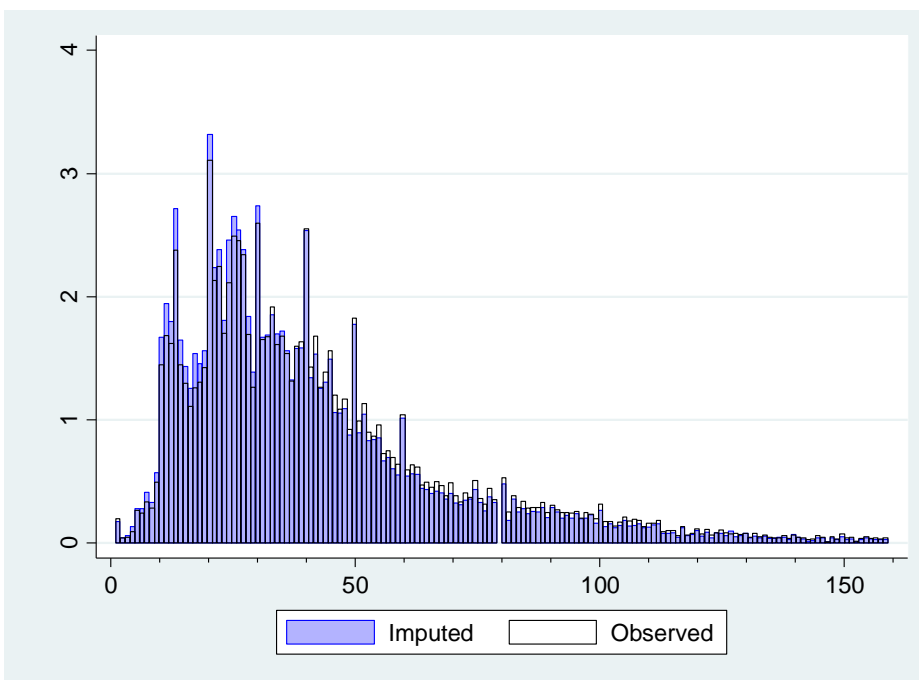


Figure A2: quantities at the pump (litres) derived from PSM using model 1

Match: Times used	Original sample			Bootstrap sample		
	Freq.	Percent	Cum.	Freq.	Percent	Cum.
1	6,239	61.32	61.32	6,102	60.39	60.39
2	2,266	22.27	83.59	2,323	22.99	83.37
3	901	8.86	92.44	899	8.9	92.27
4	373	3.67	96.11	380	3.76	96.03
5	173	1.7	97.81	206	2.04	98.07
6	97	0.95	98.76	80	0.79	98.86
7	46	0.45	99.21	44	0.44	99.3
8	32	0.31	99.53	32	0.32	99.61
9	25	0.25	99.77	14	0.14	99.75
10	7	0.07	99.84	7	0.07	99.82
11	4	0.04	99.88	7	0.07	99.89
12	8	0.08	99.96	6	0.06	99.95
13	2	0.02	99.98	4	0.04	99.99
14	2	0.02	100	0	0	99.99
15	0	0	100	1	.01	100

Table A1. Comparison of the frequency distribution of matches between PSM on the original sample and a bootstrap sample, with $e \sim N(0, 1/30625)$

Statistic	Mean of Monte Carlo samples {A} IP, PSM	SD of Monte Carlo samples {A} IP, PSM	SD of bootstrap samples e, IP, PSM	Mean of Monte Carlo samples {B}	SD of Monte Carlo samples {B}	SD of bootstrap samples
p1	1.47	0.048	0.046	1.53	0.028	0.028
p10	4.23	0.066	0.063	4.37	0.035	0.036
p25	7.82	0.099	0.096	8.02	0.050	0.053
p50	15.51	0.187	0.179	15.75	0.094	0.095
p75	30.75	0.378	0.379	30.91	0.199	0.204
p90	56.96	0.882	0.856	56.75	0.467	0.470
p99	164.89	5.193	5.259	161.20	2.918	2.940
mean	25.97	0.311	0.315	25.97	0.162	0.165
sd	35.03	2.200	1.417	34.04	0.897	0.780
skewness	6.37	3.110	1.117	6.11	1.833	0.788
N (reps)	1000	1000	1000 per MC sample	1000	1000	1000 per MC sample

Table A2. Monte Carlo (MC) simulation results

Notes

1. In each MC simulation sample we draw 42,700 observations g from a lognormal distribution selected to roughly resemble our model 2 estimates: $g \sim 15.75 \cdot e^{N(0,1)}$
2. In columns 1-3, each MC sample in sample set {A} is modified to mimic infrequent purchase (IP), replacing g with zeros according to an indicator variable (I), and inflating non-zero z values accordingly. $p(I=1)$ depends on one covariate included in the simulated dataset. We estimate $p(I=1)$ and conduct PSM to recover g . For each MC sample we also draw bootstrap samples. For each bootstrap sample we estimate $p(I=1)$, add e , conduct PSM and multiply values by \hat{ps} , as in the main text.
3. In columns 4-6 for each MC sample in sample set {B} we simply draw bootstrap samples. Comparing columns 1-3 and 4-6, it appears that the underestimated bootstrap standard errors for sd and skewness in column 3 are at least partly associated with skewness in the underlying distribution.

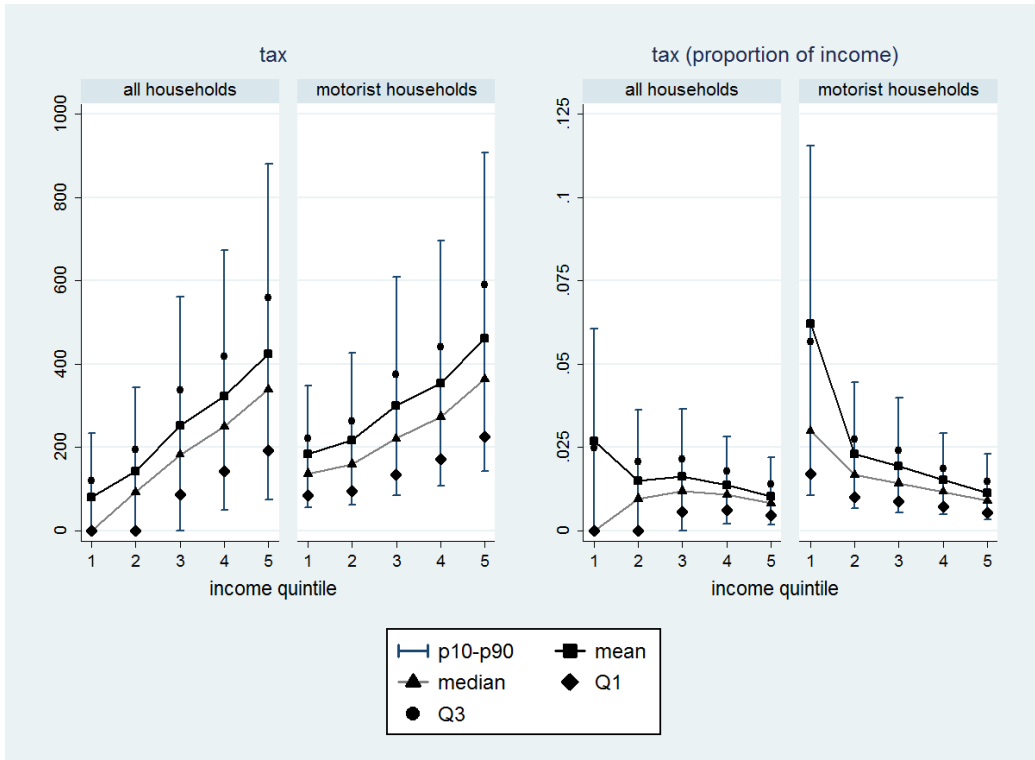


Figure A3. Estimated monetary effects of a carbon tax or ration at £100/tCO₂; model 1 ps
 Note: 'diary sample' weights applied, recalculated for the pooled sample.




Figure A4. Estimated monetary effects of 'cap and share' or 'tax and dividend' at 100/tCO₂; model 1 ps
 Note: 'diary sample' weights applied, recalculated for the pooled sample.