HEDG

HEALTH, ECONOMETRICS AND DATA GROUP

WP 19/20

Estimating and Decomposing Conditional Average Treatment Effects: The Smoking Ban in England

Matthew Robson; Tim Doran & Richard Cookson

September 2019

http://www.york.ac.uk/economics/postgrad/herc/hedg/wps/

Estimating and Decomposing Conditional Average Treatment Effects: The Smoking Ban in England

Matthew Robson, Tim Doran & Richard Cookson[‡]

Working Paper
9th September 2019

Abstract

We develop a practical method for estimating and decomposing conditional average treatment effects using locally-weighted regressions. We illustrate with an application to the smoking ban in England using a regression discontinuity design, based on Health Survey for England data. We estimate average treatment effects conditional on socioeconomic status and decompose these effects by smoking location. Results show, the ban had no effect on the level of active smoking, but significantly reduced average exposure to second-hand smoke among non-smokers by 1.38 hours per week. Our method reveals a complex relationship between socioeconomic status and the effect on passive smoking. Decomposition analysis shows that these effects stem primarily from exposure reductions in pubs, but also from workplace exposure reductions for high socioeconomic status individuals.

Keywords: Health Inequality, Equity, Conditional Average Treatment Effects, Regression Discontinuity, Heterogeneity, Smoking Ban, `lwcate`.

JEL Codes: C14, C21, C87, D63, I14, I38.

[†]EQUIPOL, Department of Health Sciences and Centre for Health Economics, University of York. Corresponding Author: Matthew Robson. Email: matthew.robson@york.ac.uk. Website: www.mrobson92.com.

[‡]We would like to thank Martin Bland, Thomas Cornelissen, Anthony Culyer, Nils Gutacker, Yuanyuan Gu, Andrew Jones, Mona Kanaan, Noemi Kreif, Andrew Mirelman, Gloria Moroni, Cheti Nicoletti, Mathilde Peron, Nigel Rice and Emma Tominey for valuable feedback and discussion. Thanks also to seminar participants at CHDS, Harvard University; CHEME, Bangor University; and HEDG, University of York, alongside conference attendees at iHEA, Basel; Economics for the Social Good, De Montford University; and the Research Impact Conference, University of York, for their helpful comments and enlightening debate. The study was funded by the Wellcome Trust (205427/Z/16/Z). All errors remain our own.

1 Introduction

Researchers often estimate average treatment effects when evaluating policies. While useful, such analysis does not provide information about how treatment effects vary for individuals with different characteristics. This makes it hard to assess the impact of policies on inequalities. To overcome this, methods that estimate treatment effects conditional on characteristics of interest are needed. Existing methods can, however, be imprecise, inflexible or inaccessible. This paper aims to provide a precise and flexible method that allows for the disaggregation of average treatment effects while maintaining accessibility and ease-of-use. Using this, the impact of policies on inequalities in society can be more easily estimated and better understood.

We therefore develop a practical econometric method for estimating and decomposing conditional average treatment effects using locally-weighted regressions. This allows complex relationships between treatment effects and conditioning variables to be estimated and the potential reasons behind these relationships to be explored. To allow others to use our method we provide a flexible Stata program: `lwcate`. We illustrate our method with an application to the smoking ban in England. We estimate the causal effects of the ban on both active and passive smoking using a regression discontinuity design based on individual-level data from the Health Survey for England. Treatment effects are then estimated conditional on socioeconomic status and decomposed by detailed smoking location variables.

An extensive literature focuses on the estimation of causal treatment effects. Rubin (1974) defines causal effects at the individual level as the difference in the *potential outcomes* an individual would have had with and without treatment. If researchers knew these individual casual effects, they could both identify the distribution of individual effects and aggregate these effects to estimate 'typical' causal effects on the population. It is, however, not possible to simultaneously observe the potential outcomes of an individual both with and without treatment. Therefore, researchers often use experiments, natural experiments or quasi-experiments where individuals are randomly or "as-if" randomly assigned to control and treatment groups.[1] By taking the difference in expected outcomes between these groups the Average Treatment Effect (ATE) can be estimated.[2]

[1]Experiments are where the researcher intentionally to groups while assignment in natural or quasi-experiments is either unintentionally randomised or non-randomised. Examples of the former include randomised controlled trails and incentivised laboratory experiments, while regression discontinuities, difference-in-differences and interrupted time-series are examples of the later two.

[2]Much literature surrounding ATEs addresses issues of compliance and selection into treatment. With distinct effects such as: the Average Treatment Effect of the Treated (ATT), Local Average Treatment

The ATE provides a useful summary of the 'typical' effect on the population. However, by focusing solely on the ATE researchers omit considerations of the distribution of individual treatment effects. The impact that any heterogeneity in individual effects has upon existing inequalities would, therefore, not be accounted for. Much other literature is, however, concerned with estimating heterogeneity in casual effects.[3] By estimating causal effects conditional on observable characteristics of interest, $\widetilde{\mathbf{x}}$, heterogeneity of treatment effects across levels of $\widetilde{\mathbf{x}}$ is allowed for. In particular, Conditional Average Treatment Effects (CATEs) estimate the expected effect of the treatment on the outcome for individuals with a specific level of $\widetilde{\mathbf{x}}$.[4] This allows researchers to identify the effects of policies on inequalities in the outcome, where the outcome is associated with $\widetilde{\mathbf{x}}$.

Standard methods for estimating CATEs involve stratification or interactions. If characteristics are discrete (e.g. sex) stratification is effective. But if characteristics are continuous (e.g. income or age) they are often either made discrete, and estimated using stratification, or are (parametrically) interacted with the treatment effect. For stratification, the setting of bounds is often arbitrary and too few discrete categories lead to bias, while too many lead to imprecision. For interactions, if the order of the parametric interaction assumed is too low the method is inflexible, leading to bias.

More nuanced methods of estimating CATEs (when $\widetilde{\mathbf{x}}$ is continuous) often use semi-parametric and non-parametric models (see Abrevaya, Hsu, and Lieli (2015) and S. Lee, Okui, and Whang (2017)). These methods, using kernel (Nadaraya, 1964; Watson, 1964) or local-linear regressions (Cleveland, 1979; Cleveland and Devlin, 1988), make no assumptions about the parametric form of the relationship between the CATEs and $\widetilde{\mathbf{x}}$. This allows for a smooth and flexible estimation of the potentially complex relation between CATEs and $\widetilde{\mathbf{x}}$, while maintaining precision.

Our proposed method similarly provides non-parametric estimates of CATEs, but is distinct from the previous literature as to do so it estimates locally-weighted regressions conditional on the quantile(s) of the observable variable(s) of interest, $q(\widetilde{\mathbf{x}})$. It provides

Effects (LATEs) and Marginal Treatment Effects (MTEs) (see J. J. Heckman and Robb Jr (1985); Imbens and Angrist (1994); Björklund and Moffitt (1987) and Cornelissen et al. (2016) for a review). In our setting these concerns over selection, and selection due to unobservables, are not of central importance. The sharp regression discontinuity design used to estimate the causal effects of the ban assumes there is no unobservable selection into treatment and that all individuals are subject to the ban.

[3]See Bedoya et al. (2017) for an extensive review of methods and a practical toolkit.

[4]Related literature which refers to CATEs are Hahn (1998); J. Heckman, Ichimura, and P. E. Todd (1997); J. Heckman, Ichimura, and P. Todd (1998); and Khan and Tamer (2010). MaCurdy, Chen, and Hong (2011) discuss the identification of CATEs, while Chang, S. Lee, and Whang (2015) and Hsu (2017) develop formal tests for CATEs.

coefficient estimates $\hat{\boldsymbol{\beta}}(\widetilde{\mathbf{x}})$, for all $\mathbf{X}$ variables within the regression model, conditional on $\widetilde{\mathbf{x}}$; of which the CATE, $\hat{\beta}\tau(\widetilde{\mathbf{x}})$, is one coefficient. The general form allows for kernel, local-linear and local-polynomial models, which use weights from either Beta, Uniform, Normal, Triangular or Epanechnikov distributions.[5] This approach is primarily developed for experiments, natural experiments and quasi-experiments, where additional $\mathbf{X}$ variables may need to be included to enable the estimation of $\hat{\beta}_{\tau}(\widetilde{\mathbf{x}})$. We provide a flexible and easy-to-use Stata program (`lwcate`) to enable other researchers to use our method. The command estimates individual-level coefficients, plots CATEs and predicts potential outcomes.

In addition to this, we conduct decomposition analysis to provide insight into the reasons underlying the estimated causal effects. Decomposition analysis is a strand of literature where differences between groups are broken into component parts. Stemming from two seminal papers (Blinder, 1973; Oaxaca, 1973) the Oaxaca-Blinder decomposition identifies the *endowment*, *coefficient* and *unexplained* components, in order to understand average differences better. When the two groups concerned are control and treated groups, as Słoczyński (2015) shows, the average treatment effect can be similarly decomposed. We follow this approach but apply it to the regression discontinuity framework.

Taking inspiration from Machado and Mata (2005); Melly (2005); and Chernozhukov, Fernández-Val, and Melly (2013), who decompose differences conditional on the quantiles or levels of the *outcome* variable, we decompose CATEs, which are conditional on the level of an *explanatory* variable. This allows the heterogeneity in CATEs to be better understood, as the underlying components of these effects are also estimated conditionally. This sheds light on the reasons behind any observed heterogeneity. Knowing these reasons could prove useful for future policy: allowing for the prevention of further unintentional heterogeneity.

We apply our methods to the smoking ban, which was introduced by the UK government across England in July 2007. The ban effectively prohibited the smoking of tobacco in enclosed public places, including workplaces and places where the public accessed goods and services, including public transport, private clubs and pubs. Exemptions were allowed for psychiatric units, nursing homes, prisons and other specified public places. Those who breached the ban were subject to fines, up to £200 for individuals and £2,500 for businesses.

Early evaluations of the smoking ban found that it was successful in reducing exposure to second-hand smoke in public places but had limited impact on tobacco consumption,

[5]While the Uniform, Normal, Triangular or Epanechnikov are standard in the literature, we propose and formulate the Beta distribution as an alternative. It is particularly suited to our quantile approach and avoids issues of bias at the extremes.

suggesting that smokers changed the locations where they smoked but not the amount. Evidence shows that quit rates for smokers were increasing prior to the ban but remained low, at less than 3% per year. Interviews with high cardiovascular risk smokers in Scotland suggest introduction of the ban was associated with a small short-term increase in quit rates but this was not sustained (Fowkes et al., 2008). Difference-in-differences studies also found very limited short-run effects on smoking prevalence and total level of smoking (J. Lee, Glantz, and Millett, 2011; Jones et al., 2015). In contrast, smokers reported substantial drops in smoking at work (from 14% to 2%) and in pubs (from 34% to 2%), locations covered by the primary legislation, but also inside the home (from 65% to 55%), which was not covered (J. Lee, Glantz, and Millett, 2011). This contrasts with evidence from the US, where reduced exposure in public places following smoking bans was accompanied by increased exposure in private areas due to substitution (Adda and Cornaglia, 2010).

The evidence on the equity impacts of smoking bans is less clear. Because smoking prevalence and second-hand exposure is higher in more deprived groups, legislation which restricts smoking has the potential to reduce socioeconomic inequalities in health. However, Sims et al. (2011) show that although the ban reduced second-hand smoke exposure amongst non-smokers, significant reductions were not found for lower social class households. A review of the impacts on socioeconomic inequalities in health of national smoking strategies found only six studies based on comprehensive national/state smoke free legislation; two reported equity neutral results and four reported a negative impact (Brown, Platt, and Amos, 2014). All UK studies reported overall reductions in second-hand exposure for children but in Wales and Northern Ireland this was concentrated in higher socioeconomic groups (G. Moore, Holliday, and L. Moore, 2011; G. Moore, Currie, et al., 2012; Jarvis et al., 2012) In Scotland higher absolute reductions for lower socioeconomic groups still resulted in relative increases in inequalities between socioeconomic groups (Akhtar et al., 2010).

This paper contributes to the methodological literature surrounding the estimation of causal treatment effects and to applied research on smoking and health inequalities. We develop a novel econometric method to estimate and decompose conditional average treatment effects using locally-weighted regressions and provide a flexible, easy to use, program in Stata. We further estimate causal effects for both *active* and *passive* smoking in England; estimating and decomposing these treatment effects, conditional on socioeconomic status.

The remainder of the paper is organised in six main sections. *Methods* formulates the potential outcomes framework (2.1), regression discontinuity design (2.2), conditional average treatment effects estimator (2.3) and decomposition analysis (2.4). *Data* describes

the survey (3.1), explains the derivation of the socioeconomic status variable (3.2), shows descriptive statistics (3.3) and plots smoking outcomes over time (3.4). Then, *Results* show estimates for average treatment effects (4.1) and conditional average treatment effects (4.2). Active smoking behaviour (4.3) and passive smoking exposure (4.4) are then investigated in more detail, providing counterfactual predictions, decomposition analysis and conditional decomposition analysis. Following that are the *Discussion*, *Conclusion* and *Appendix*.

2 Methods

2.1 Potential Outcomes Framework

To provide cohesion and an overarching framework to the below methods, the potential outcomes framework is used. For a population of N individuals we observe *actual* outcomes $\mathbf{y}$. We also observe $\mathbf{D}$ which shows if an individual *is* in the control group ($D = 0$) or in the treatment group ($D = 1$). We can imagine the *potential* outcomes an individual would have if an individual were in the control group, with $\mathbf{Y}(0)$, or in the treatment group, with $\mathbf{Y}(1)$. It is D which determines whether the potential outcome for an individual is *observable* or *counterfactual* (i.e. not observable). If $D = 0$ then $Y(0)$ is observable and $Y(1)$ is counterfactual, and vice versa for $D = 1$. The outcome we observe, $\mathbf{y}$, can therefore be written as follows:

$$\mathbf{y} = \mathbf{Y}(0)(1 - \mathbf{D}) + \mathbf{Y}(1)\mathbf{D} \tag{1}$$

The effect of receiving the treatment, for one individual, would be the difference in their potential outcomes in the control and treatment group: $\tau_I = Y(1) - Y(0)$. These individual causal effects can be aggregated to establish the 'typical' effects of the treatment. The Average Treatment Effect (ATE), τ_A, is the expected difference in potential outcomes, across the whole sample, while the Conditional Average Treatment Effect (CATE), $\tau_C(\widetilde{\mathbf{x}})$, is the expected difference in potential outcomes for a particular sub-sample, identified by $\widetilde{\mathbf{x}}$:

$$\tau_A = E[\boldsymbol{\tau}_I] = E[\mathbf{Y}(1) - \mathbf{Y}(0)] \tag{2}$$

$$\tau_C(\widetilde{\mathbf{x}}) = E[\boldsymbol{\tau}_I | \widetilde{\mathbf{X}} = \widetilde{\mathbf{x}}] = E[\mathbf{Y}(1) - \mathbf{Y}(0) | \widetilde{\mathbf{X}} = \widetilde{\mathbf{x}}] \tag{3}$$

Just as the 'typical' effects can be seen as the aggregation of individual effects, causal effects can be seen as the aggregation of separate *components*. For example, let the outcome,

$\mathbf{y}$ be the number of cigarettes smoked. This could be split into two components, the number smoked at home (A) and the number smoked elsewhere (B), as shown below. If a policy aimed at reducing the number of cigarettes smoked, it might be of interest to decompose the 'total' treatment effect into two such component parts. The causal effect of the policy in reducing the number of cigarettes smoked at home (Component A) and elsewhere (Component B). This decomposition could be generalised to any number of components and elicited at either the individual, average or conditional average level. Indeed, if individual or conditional causal effects were heterogeneous, this decomposition may shed light onto the underlying reasons for such heterogeneity.

$$\mathbf{y} = (\mathbf{Y}_A(0) + \mathbf{Y}_B(0))(1 - \mathbf{D}) + (\mathbf{Y}_A(1) + \mathbf{Y}_B(1))\mathbf{D} \tag{4}$$

$$\tau_I = \underbrace{(Y_A(1) - Y_A(0))}_{Component\ A} + \underbrace{(Y_B(1) - Y_B(0))}_{Component\ B} \tag{5}$$

Unfortunately, it is not possible to *observe* counterfactual outcomes, therefore treatment effects and counterfactual outcomes need to be *estimated*. The following sections present the estimation procedures for the above. Section 2.2 shows the regression discontinuity design used to estimate the ATE, τ_A. Section 2.3 proposed a method to non-parametrically estimate CATEs, $\tau_C(\widetilde{\mathbf{x}})$. Section 2.4 formulates the method to decompose ATEs, then Section 2.5 extends this to decompose CATEs. Through each of these methods not only are the causal effects estimated, but the counterfactual outcomes can be predicted for each individual. Within the ATE methods homogeneous treatment effects are assumed, but the CATE methods allow for heterogeneity, conditional on some $\widetilde{\mathbf{x}}$, aiming to get one step closer to estimating the individual level causal effects and predicting counterfactual outcomes.

2.2 Regression Discontinuity

In order to estimate the causal effect of the smoking ban on our outcomes of interest we use the Regression Discontinuity Design (RDD) originally proposed by Thistlethwaite and Campbell (1960). An RDD is appropriate when a single continuous forcing variable is used to determine whether an individual is in the control or treatment group. Individuals above a 'threshold' are assigned to the treatment group, while those below are assigned to the control. Because of this 'as if random' assignment, the expected characteristics of individuals within

each group should be equal. Therefore, regressions can be used in neighbourhoods close to the threshold to identify the causal effect of the treatment on outcomes on interest.

Within our analysis the forcing variable used is time (i.e. the interview date) while the threshold is determined by the time of the ban (1st July 2007). If the interview date of the survey is after the date of the ban individuals are classed as treated, while those interviewed before are in the control group.[6] We utilise a *sharp* RDD as the ban was introduced nationally on the 1st July 2007.[7] Our general econometric model is as follows:

$$\mathbf{y} = \beta_0 + \mathbf{D}\beta_1 + \mathbf{Z}\beta_2 + (\mathbf{D}' \times \mathbf{Z})\beta_3 + \mathbf{u}, \qquad for\ |\mathbf{Z}| \leq h \tag{6}$$

Where, for N observations, $\mathbf{y}$, the outcome of interest, $\mathbf{D}$, treatment dummy, $\mathbf{Z}$, forcing variable (days from the ban), and $\mathbf{u}$, error term, are $[N \times 1]$ vectors. $\mathbf{Z}$ is centred so that $\mathbf{Z} = 0$ at the point of the intervention; ensuring ease of interpretation of the coefficients. The bandwidth, h, is used to define the size of the neighbourhood surrounding the threshold. Coefficient β_0 denotes the expected level of $\mathbf{y}$ for the control group when $\mathbf{Z} \to 0$. The main coefficient of interest is β_1 which estimates the Average Treatment Effect, $\hat{\tau}_A$, by identifying the size of the discontinuity at the threshold. Coefficients β_2 and β_3 denote the trend of $\mathbf{y}$ over time and the change in that trend due to the treatment, respectively.

2.3 Conditional Average Treatment Effects

Average Treatment Effects, τ_A, are useful at summarising the *typical* effect on the whole population. However, they ignore how treatment effects can differ for individuals with different characteristics (e.g. SES, age, gender). Conditional Average Treatment Effects, $\tau_C(\widetilde{\mathbf{x}})$, can be estimated to allow for such heterogeneity, conditional on observable characteristics $\widetilde{\mathbf{x}}$.

We propose a smooth and flexible non-parametric method for estimating treatment effects, conditional on $\widetilde{\mathbf{x}}$ the observable characteristic(s) of interest. The approach estimates a

[6] This means our design is somewhat similar to an Interrupted Time Series (ITS) design. However, unlike ITS in general, our data are cross-sectional, we use a relatively small time period surrounding the intervention and utilise alternative bandwidths to determine the size of the neighbourhood as sensitivity. Using time is subject to two major criticisms: 1) other factors could also change at the same time as the ban; 2) enforcement of the ban and/or behaviour change in response to it may be anticipatory or delayed. Sensitivity analysis will strive to address this.

[7] RDDs are defined as either "sharp" or "fuzzy". For the former, the probability of belonging to the treatment group is deterministically determined by the forcing variable, while for the latter it is only an increase in the probability of belonging to the either group when moving away from the threshold which is required. Fuzzy RDD's rely on using the assignment to treatment as an instrumental variable, to identify the effect of the treatment on the 'compliers'.

set of locally-weighted regressions; where the 'local' weights are a function of the rank of $\widetilde{\mathbf{x}}$, rather than the level of $\widetilde{\mathbf{x}}$. The approach allows for the estimation of Conditional Average Treatment Effects, $\hat{\tau}_C(\widetilde{\mathbf{x}}) = \hat{\beta}_\tau(\widetilde{\mathbf{x}})$, and other conditional coefficients, $\hat{\boldsymbol{\beta}}(\widetilde{\mathbf{x}})$ (for additional observable variables X), alongside the prediction of potential outcomes, $Y(0)$ and $Y(1)$. In general, we model:

$$\mathbf{y} = \beta_0(\widetilde{\mathbf{x}}) + \beta_\tau(\widetilde{\mathbf{x}})\mathbf{D} + \boldsymbol{\beta}(\widetilde{\mathbf{x}})\mathbf{X} + \mathbf{u} \tag{7}$$

Where $\mathbf{y}$ is the outcome variable, $\mathbf{D}$ the treatment dummy, $\mathbf{X}$ the additional explanatory variables and $\mathbf{u}$ the error term. The coefficients are functions of the observable characteristic(s) of interest $\widetilde{\mathbf{x}}$: the constant, $\beta_0(\widetilde{\mathbf{x}})$, CATE, $\beta_\tau(\widetilde{\mathbf{x}})$, and additional coefficients, $\boldsymbol{\beta}(\widetilde{\mathbf{x}})$.

To obtain estimates of the above model our approach runs a set of M locally-weighted regressions. Where in each $j \in M$ regression, weights, w_{ij}, are derived from a kernel function; such that $\sum_i^N (w_{ij} r_i) / \sum_i^N (w_{ij}) \approx q_j^*$, the weighted average of the rank of the individual equals the quantile of interest, $\forall j \in M$. Where for individuals $i \in N$, let $r_i = r(\widetilde{x}_i)$ be the expected normalized rank of individual i with respect to $\widetilde{x}_i$, $q_j^* = q_j^*(\widetilde{x})$ be the expected quantile of interest for local-regression j and w_{ij} be the weight assigned to i in regression j.[8]

Once estimated, the set of M regressions provide separate coefficients $\hat{\beta}_0(\widetilde{\mathbf{x}}_i), \hat{\beta}_\tau(\widetilde{\mathbf{x}}_i)$ and $\hat{\boldsymbol{\beta}}(\widetilde{\mathbf{x}}_i)$ for individuals, i, belonging to quantile j. These coefficients are estimated parametrically, within each j regression, but non-parametrically, across the M quantiles of interest. As an example, for an RDD, the model could be:

$$\mathbf{y} = \beta_0(\widetilde{\mathbf{x}}) + \beta_\tau(\widetilde{\mathbf{x}})\mathbf{D} + \beta_2(\widetilde{\mathbf{x}})\mathbf{Z} + \beta_3(\widetilde{\mathbf{x}})(\mathbf{D}' \times \mathbf{Z}) + \mathbf{u}, \qquad if |\mathbf{Z}| < h \tag{8}$$

The approach can be estimated as a kernel, local-linear or local-polynomial regression; these alternatives are explicitly formulated within Appendix A.1. Appendix A.2 provides extensive simulation tests of the proposed model; comparing the mean-squared error and precision of estimated treatment effects.

[8]Note that here we use q_j^* and r_i to denote the *middle* value of the quantile and normalised rank of $\widetilde{\mathbf{x}}$, i.e. if we have two quantiles/ranks we mean not that the first quantile/rank equals 0.5, but 0.25, and the second 0.75, not 1.

2.3.1 Weights

In order to estimate each locally-weighted regression, individual weights need to be derived. Within the literature surrounding kernel and local-linear regressions these weights are typically derived from either the Uniform, Normal, Triangular or Epanechnikov distributions; the choice of which tends not to be crucial (Pagan and Ullah, 1999). Here we propose and formulate an alternative, the Beta distribution, which is particularly suited for the quantile approach we use.

The Beta distribution is bounded between 0 and 1, using two flexible shape parameters, a and b. We make two assumptions: (1) $E(Q) = q^*$: the expected value of Q is equal to our quantile of interest, and (2) $Var(Q) = \frac{q^*(1-q^*)}{s}$: the variance of Q is dependent on a precision parameter, s (where $s \geq 3$), and decreases towards the bounds. From this, and according to the characteristics of the Beta distribution, a and b are obtained as functions of q^* and s.

$$E[Q_j] = \frac{a_j}{a_j + b_j} = q_j^* \tag{9}$$

$$Var[Q_j] = \frac{a_j b_j}{(a_j + b_j)^2(a_j + b_j + 1)} = \frac{q_j^*(1 - q_j^*)}{s} \tag{10}$$

It follows that:

$$a_j = q_j^*(s-1), \qquad b_j = (1 - q_j^*)(s-1) \tag{11}$$

With these equations, for a given s, the probability density function (PDF), and therefore the cumulative distribution function (CDF), can be derived for each level of q^*. Weights, w_{ij}, for each regression $j \in M$ can then be calculated for each individual, $i \in N$, by taking the integral of the PDF for the interval within which the individual is situated.[9] Formally:

$$w_{ij} = \int_{r_i - \frac{1}{2N}}^{r_i + \frac{1}{2N}} \left(\frac{\Gamma(a_j + b_j)}{\Gamma(a_j)\Gamma(b_j)} \left(q^{a_j - 1}(1-q)^{b_j - 1} \right) \right) dq \tag{12}$$

To illustrate the above, let there be five quantiles at which we want to estimate local regressions. In order to run five locally-weighted regressions, weights w_{ij} are needed, where $j \in 1, .., 5$ and $i \in 1, .., N$. Figure 1 plots the PDFs for $q^* = [0.1, 0.3, 0.5, 0.7, 0.9]$, respectively for each j, when $s = 20$. The weight given to each individual i for regression j is calculated

[9] For example, if $N = 10$, then for the first individual their expected normalised rank is 0.05. To obtain their individual weight, w_{ij}, take the integral of the probability density function for the interval $[0.05 - \frac{1}{2N}, 0.05 + \frac{1}{2N}]$ or $[0, 0.01]$. This is equivalent to taking the CDF(0.01) - CDF(0).

from Equation 12. The weights for a given i are generally higher when $r_i \to q_j^*$. In other words, when we are interested in identifying the CATE, at q_j^*, so generally, those who are closer to q_j^* are given a higher weight. The higher the number of quantiles (groups) we split the observations into the more regressions will be run. The higher the level of precision, s, the steeper the PDF and higher the variance in w_{ij}.

Figure 1: Beta Distribution: Probability Density Function

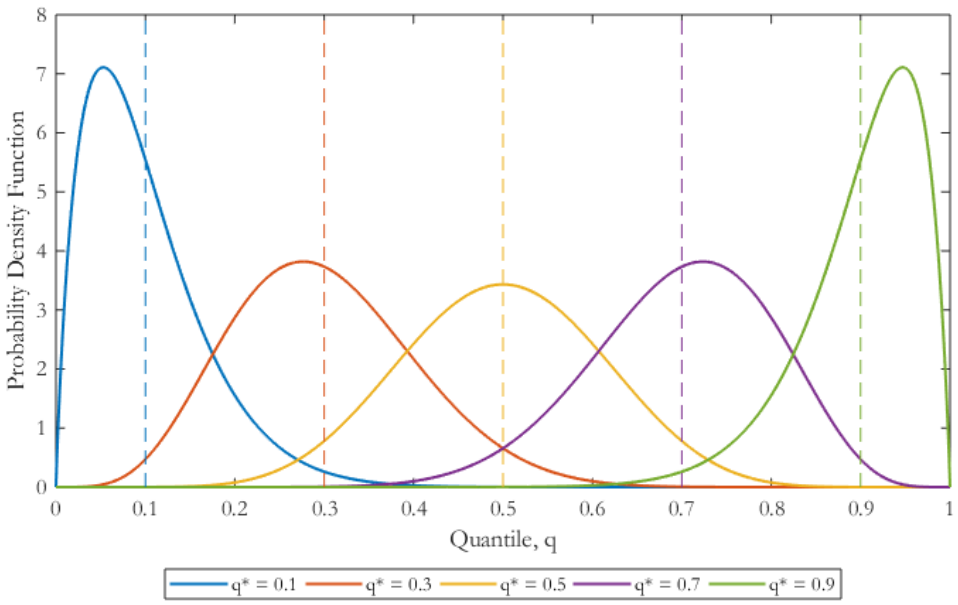

If the conditioning variable, $\widetilde{X}$, is continuous then the individual ranks, r_i are unique. However, it is not necessary for X to be continuous in order to calculate weights. For κ equal ranks, the individual weight is calculated as in Equation 12, but interval would be between the lowest 'potential' rank r_{Li} and the highest 'potential' rank r_{Hi} of that group of individuals, where $r_{Hi} = r_{Li} + \kappa$. That weight is then divided by κ to give w_{ij}.

Often, when conditioning on $\widetilde{\mathbf{X}}$ only one variable is used (e.g. SES). However, multiple variables can be conditioned upon (e.g. both SES and age). Within our method this is achieved simply by multiplying weights from marginal Beta distributions. However, an alternative the Dirichlet (a multinomial Beta) distribution is formulated in Appendix A.3. Alongside this, an alternative constant variance assumption for the Beta weights is derived.

2.3.2 Locally-Weighted Regressions

Once weights, w_{ij} (here denoted as $\mathbf{w}_j$), have been calculated the set of M weighted least-squares regressions can be run. Within each j weighted least squares regression we model:

$$\sqrt{\mathbf{w}_j}\mathbf{y}_j = \sqrt{\mathbf{w}_j}\left(\beta_{0j} + \beta_{\tau j}\mathbf{D}_j + \boldsymbol{\beta}_j\mathbf{X}_j + \mathbf{u}_j\right) \tag{13}$$

Where each j regression is multiplied by the square-root of the weights $\mathbf{w}_j$. The estimated coefficients $\hat{\beta}_{0j}, \hat{\beta}_{\tau j}$ and $\hat{\boldsymbol{\beta}}_j$ are assigned to each individual i, belonging to quantile j, providing the coefficients $\hat{\beta}_0(\widetilde{\mathbf{x}}_i), \hat{\beta}_\tau(\widetilde{\mathbf{x}}_i)$ and $\hat{\boldsymbol{\beta}}(\widetilde{\mathbf{x}}_i)$ which are conditional on i's $\widetilde{\mathbf{x}}_i$. Within Stata this is done by running the regression with *aweights* or *pweights*.[10] Hence, it can flexibly be used for many types of regression, for example, OLS, probit and tobit regressions.

2.4 Decomposition

While the above methodology allows for the estimation of treatment effects, it reveals little of the reasons why such an effect exists. Through using decomposition analysis, popularised by Blinder (1973) and Oaxaca (1973), we can estimate the component parts of the treatment effect; which may shed light on such reasons. In order to do so, we estimate two separate equations for the control and treatment group:

$$\mathbf{y}_0 = \beta_{00} + \mathbf{X}_0\boldsymbol{\beta}_0 + \mathbf{u}_0 \tag{14}$$

$$\mathbf{y}_1 = \beta_{10} + \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{u}_1 \tag{15}$$

Where the subscript, g, denotes the *control* group ($g = 0$) and *treatment* group ($g = 1$). The dependent variable $\mathbf{y}$ is an $[N \times 1]$ vector, for N observations. The p independent variables are represented by the $[N \times p]$ matrix $\mathbf{X}$, while $\mathbf{u}$ represents the normally distributed $[N \times 1]$ error vector. The intercept is denoted by β_{g0} and the $[p \times 1]$ coefficient vector $\boldsymbol{\beta}$ denotes the effects of the independent variables on the dependent variable. Given (14) and (15), we are interested in the difference between the expectation of $\mathbf{y}_0$ and $\mathbf{y}_1$, which gives the average treatment effect. Through rearranging the equations, we can decompose this effect we can estimate the components of this effect:

[10] See Dupraz (2013) for a useful summary of weighted least squares in Stata.

$$
\begin{aligned}
E[\mathbf{y}_1] - E[\mathbf{y}_0] &= E[\beta_{10} + \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{u}_1] - E[\beta_{00} + \mathbf{X}_0\boldsymbol{\beta}_0 + \mathbf{u}_0] \\
&= \beta_{10} + E[\mathbf{X}_1]\boldsymbol{\beta}_1 - (\beta_{00} + E[\mathbf{X}_0]\boldsymbol{\beta}_0) \\
&= \underbrace{E[\mathbf{X}_1](\boldsymbol{\beta}_1 - \boldsymbol{\beta}_0)}_{Coefficients} + \underbrace{(E[\mathbf{X}_1] - E[\mathbf{X}_0])\boldsymbol{\beta}_0}_{Endowments} + \underbrace{\beta_{10} - \beta_{00}}_{Unexplained}
\end{aligned} \tag{16}
$$

The *coefficients* component denotes the proportion of the treatment effect which is attributable to a change in coefficients. The *endowments* component identifies how much of the difference is attributable to differing endowments of the two groups. The *unexplained* component simply shows how much of this difference in unexplained by the previous two terms. These components can be separated for each explanatory variable within the model. Estimation is done using the *oaxaca* command in Stata (Jann, 2008).

2.4.1 Regression Discontinuity Particulars

While the above describes decomposition analysis in general, there are however some aspects particular to the regression discontinuity design. Within a regression discontinuity design we are interested in the estimating the discontinuity at the threshold (0), not the difference in the expectation of the outcome before and after the threshold in general. The treatment effect of interest is:

$$
\lim_{z \downarrow 0} E[\mathbf{y}_1] - \lim_{z \uparrow 0} E[\mathbf{y}_0] \tag{17}
$$

As a result, when specifying Model (14) and (15) the J independent variables must be interacted with a forcing variable, Z (which is centred on zero), in order to estimate the treatment effect of interest. Our matrix of independent variables $\mathbf{X}$ should be:

$$
\mathbf{X} = \begin{bmatrix}
x_{11} & \dots & x_{1J} & z_1 & z_1x_{11} & \dots & z_1x_{1J} \\
x_{21} & \dots & x_{2J} & z_2 & z_2x_{21} & \dots & z_2x_{2J} \\
\vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\
x_{N1} & \dots & x_{NJ} & z_N & z_Nx_{N1} & \dots & z_Nx_{NJ}
\end{bmatrix}
$$

With X so specified, we are interested in the components which correspond to $\beta_j, \forall j \in [1...J]$. They provide the estimates when $z = 0$ and so relate to the discontinuity at the point of the intervention. Of course, researchers interested in how trends shift due to the treatment will find the $j \in [J + 1...2J + 1]$ coefficients of interest.

Combining decomposition analysis with the regression discontinuity design provides useful additions and checks. Independent variables we expect to change, due to the treatment, could provide insights into the reasons behind the treatment effect. If *endowment* components change significantly, then (if the treatment *caused* those independent variables to change) the component of the 'total' treatment effect attributed to that change in the independent variable can be estimated. Similarly, if the *coefficient* component attributable to that independent variable changes the explanatory power that variable had, on explaining the dependent variable, would also have changed.

Observable characteristics (controls) need to be balanced across control and treatment group to adhere to the underlying assumptions of the RDD. Through checking the *endowments* component of those controls we can check balance between groups; the coefficients should all be insignificant. Furthermore, the *coefficients* component will identify any heterogeneity in treatment effects for the control variables (e.g. a larger effect for females).

The decomposition is explained in more detail in Appendix A.4, alongside an alternative two-stage model which could allow for a causal interpretation of the endowment effects.

2.5 Decomposing Conditional Average Treatment Effects

The above methods can be combined to increase the explanatory power of the analysis. By using the decomposition analysis in each of the M locally-weighted regressions we can decompose conditional average treatment effects, conditional on the observable variable of interest, $\widetilde{\mathbf{x}}$. This allows for a deeper understanding of the reasons underlying the relationship between the CATEs and $\widetilde{\mathbf{x}}$.

3 Data

3.1 Health Survey for England

The Health Survey for England provides rich cross-sectional data on an extensive range of variables relating to demographic and household variables. Particular interest is paid to health-related variables, with a follow-up nurse's visit conducted for most households. Interviews are conducted, with all eligible members of each household, for a representative sample across England. We focus on adults in the 2007 wave ($n \approx 6300$).

The outcome measures of interest relate to both *active* and *passive* smoking. For active smoking we focus on *smoker* (whether an individual currently smokes) and *packs* (the number of cigarettes smoked in a day, averaged over the week, in packs of five). Additionally, we use in-depth responses of where an individual smokes (at home, work, other's home, pubs, cafes, shops or car) and whom they smoke near (babies, children, youths, the elderly, those who are pregnant or those with ashthma). Throughout, when analysing active smoking, we restrict the sample to those individuals who are ever-smokers (i.e. those who have ever smoked in their lives ($N \approx 3299$)).

The outcome variables for passive smoking are *exposed* (whether an individual is exposed to other people's tobacco smoke) and *hours exposed* (the typical number of hours per week an individual is exposed to other people's tobacco smoke).[11] More detailed information is also available for where individuals are exposed to second-hand smoke (at home, work, other's home, public transport and pubs). For analysis of passive smoking the sample is restricted to non-smokers (i.e. those who are currently not smoking ($N \approx 4693$)).[12]

In addition to the outcome variables there are additional demographic and household characteristics. We use age, sex, married, ethnicity (white = 1), urban/rural, pregnant and ever-smoked. Detailed information on variables relevant to socioeconomic status (SES) is available. These variables include: log equalised income[13], Index of Multiple Deprivation (IMD) deciles, savings and investments, benefit recipient, NS-SEC occupational classifications, labour market status, highest qualifications, housing tenure and overcrowing. A multivariate index for SES is calculated using these variables in Section 3.2; it is this which we use in our main analysis. The time of the interview is particularly important for our analysis which is available as the month of interview. For conciseness, and relevance to the regression discontinuity design used, descriptive statistics of the variables used are shown in Section 3.3.

[11]The wording of the questionnaire for the main outcome variables is: *smoker* "Do you smoke cigarettes at all nowadays?"; *packs* "how many cigarettes a day do you usually smoke on weekdays[/weekends]?"; *hours exposed* "Now, in most weeks, how many hours are you exposed to other people's tobacco smoke?"; *exposed* is a dummy variable derived from hours exposed. For smokers the where and near whom variables relate to the last seven days, while for non-smokers the where exposed variables ask if they are "regularly exposed".

[12]In addition to self-reported outcomes relating to smoking, cotinine levels are measured in a follow up nurses visit. Cotinine, derived from saliva samples, is a metabolite of nicotine and is commonly used as a biomarker of tobacco smoke exposure. While not used in the main analysis it is used in Appendix A.5, giving confidence in the results which use self-reported measures of both active and passive smoking.

[13]See Appendix A.6 for issues relating to missing observations and details on imputation to address this.

3.2 Socioeconomic Status

Socioeconomic status (SES) is a complex construct consisting of the accumulation and interaction of income, wealth, education, occupation, geographic region, housing, social networks and culture, amongst other factors. Yet, researchers often focus on singular dimensions of socioeconomic status. In doing so the socioeconomic status of individuals (or households) may be misclassified and the impact that other dimensions have upon the question of interest thereby missed.

Our approach aims to use a data driven approach to address the above concerns and the multidimensional nature of SES. We use principal component analysis (PCA) to estimate the latent construct of SES.[14] By using the weights from the first component of the PCA we estimate the relation between each individual factor and overall SES. The higher the (positive) weight, the more that factor denotes an increase in SES. It is the combination and interaction of an array of factors which will determine each individual's SES. This method derives a continuous numerical index, which determines the relative SES of an individual within the sample. The index embodies both continuous variables (e.g income) and more discrete class distinctions (e.g. NS-SEC). It incorporates the more fluid variables, such as income and labour market status, alongside more crystallised variables, such as housing and qualifications. Unfortunately, variables such as parental class or social and cultural capital are not available in the survey. However, the methodology taken is pragmatic and flexible and allows for important variables to be included if and when they are available from different datasets.

We estimate SES conditional upon age and identify the relative socioeconomic status of an individual within their age category. This separation allows factors to have different weights throughout the lifecourse. This index, therefore, more closely resembles life-time potential SES. The weights for each factor, conditional on the age of the individual, are shown in Table 1.

Across age categories we observe consistent patterns. Having a higher income, living in a less deprived area (IMD), possessing savings, having higher level occupational jobs, owning your house, and having higher qualifications give a higher level of SES. While relying on benefits, never having been employed, currently being unemployed, having low levels of qual-

[14]This method draws from the development literature, where asset-wealth indices are similarlty constructed (see Vyas and Kumaranayake (2006)).

Table 1: Socioeconomic Status Index, Conditional on Age; PCA

	18-24 Weights	**25-34** Weights	**35-44** Weights	**45-54** Weights	**55-64** Weights	**65+** Weights
Economic						
Log Equivalised Income	0.362	0.385	0.399	0.375	0.377	0.339
IMD Quintile	0.213	0.202	0.250	0.273	0.281	0.299
Savings & Investments	0.157	0.142	0.157	0.176	0.235	0.256
No Benefits	0.298	0.287	0.237	0.248	0.302	0.288
NS-SEC						
- Higher Manager/Profess.	0.076	0.180	0.196	0.163	0.158	0.172
- Lower Manager/Profess.	0.113	0.194	0.162	0.193	0.210	0.217
- Intermediate	0.148	-0.006	-0.003	0.028	0.018	0.066
- Small Employ./Own Acc.	0.005	-0.016	-0.024	0.020	0.032	0.031
- Lower Superv./Tech.	0.033	-0.015	-0.017	-0.022	-0.058	-0.056
- Semi-routine Occup.	-0.028	-0.163	-0.120	-0.137	-0.136	-0.140
- Routine Occup.	-0.017	-0.123	-0.167	-0.238	-0.214	-0.212
- Never Employed	-0.274	-0.155	-0.189	-0.132	-0.059	-0.042
Labour Market Status						
- Employed	0.306	0.321	0.316	0.319	0.267	0.071
- Unemployed	-0.137	-0.086	-0.078	-0.068	-0.071	
- Retired			-0.010	0.007	-0.059	-0.046
- Full Time Student	-0.021	-0.038	-0.030	-0.043	0.002	
- Looking After Home	-0.289	-0.242	-0.228	-0.154	-0.101	-0.008
- Other Inactive	-0.114	-0.163	-0.186	-0.266	-0.228	
Highest Qualification						
- Degree	0.171	0.238	0.227	0.211	0.239	0.192
- Higher Education	0.055	0.053	0.047	0.077	0.107	0.142
- A Level	0.130	0.014	0.046	0.053	0.065	0.085
- GCSE A*-C	-0.081	-0.087	-0.082	-0.011	0.065	0.145
- GCSE D-G	-0.066	-0.114	-0.069	-0.048	-0.028	0.003
- Other	-0.035	-0.014	-0.041	-0.033	-0.002	0.057
- None	-0.254	-0.229	-0.242	-0.279	-0.327	-0.336
Housing Tenure						
- Own Outright	0.082	0.020	-0.007	-0.009	0.174	0.347
- Own - Mortgage/Loan	0.297	0.295	0.300	0.271	0.055	-0.043
- Part Rent/Part Mort.	-0.018	0.004	-0.023	-0.014	0.018	-0.040
- Rent	-0.345	-0.313	-0.324	-0.322	-0.302	-0.350
- Live Rent Free	0.018	0.015	-0.014	-0.008	0.025	-0.021
Overcrowding	-0.205	-0.239	-0.191	-0.139	-0.197	-0.199
Observations	452	886	1156	1039	995	1475

ifications, renting accommodation and living in overcrowded accommodation all correspond to lower SES. These closely correspond to our prior intuition.

By estimating SES conditional on age, we also observe interesting differences between age groups. Income becomes more important in mid-life, while living in a less deprived area and having savings becomes increasingly important as individuals get older. Being employed in higher-level occupations is important at older ages, whereas for those who are younger (and perhaps more transient workers) this is less crucial. Currently being employed is crucial for those of working age, but for those at retirement age it has little importance. Owning a house is important throughout, but for those younger than 54 the majority own their house

through a mortgage (85.2% of home owners), so this is the signal of higher SES; however if individuals are approaching or have reached retirement age (55+) it is important that they have paid off the mortgage (only 7.9% of homeowners older than 65 have still have a mortgage/loan).

As with all measures of SES, our construct has limitations, but for our purposes it represents a flexible data-driven attempt to move closer to the multidimensional construct that socioeconomic status is.[15]

3.3 Descriptive Statistics

The descriptive statistics for variables relating to smoking outcomes and demographic characteristics are shown in Table 2. The first four columns show the mean values and number of observations for each variable before and after the ban, while the final column shows the difference in means and significance levels between these two groups from a two-sample t-test. The data are cross-sectional, with a sample of approximately 6,300 adults from the 2007 wave. The differences in the smoking outcomes provide an overview and hint at the later results, while the lack of difference in demographic characteristics highlights the balanced nature of the survey before and after the ban.

Table 2: Descriptive Statistics

	Before		**After**		**Difference**
	Mean	Obs	Mean	Obs	
Smoking Outcomes					
Smoker	0.21	3220	0.22	3080	0.01
Packs Smoked Per Day	0.54	3223	0.53	3079	-0.02
Exposed to Smoke	0.48	3205	0.26	3072	-0.22***
Hours Exposed to Smoke	5.28	3205	3.15	3072	-2.26***
Ever Smoked	0.60	3221	0.61	3081	0.00
Cotinine Levels	68.92	2104	59.69	1777	-9.38**
Smoke in Pubs	0.07	3204	0.01	3065	-0.07***
Exposed in Pubs	0.26	2926	0.02	2889	-0.24***
Demographics					
Age	49.73	3237	50.45	3094	0.43
Sex (1 = female)	0.56	3237	0.55	3094	-0.00
SES	0.59	3048	0.60	2955	0.01
Married Couple	0.57	3237	0.56	3094	-0.01
Ethnicity (1 = white)	0.89	3222	0.91	3086	0.02**
Urban	0.78	3237	0.77	3094	-0.01
Pregnant	0.01	3237	0.01	3094	0.00

[15]For further detail, Appendix A.7 cross-tabulates smoking outcomes, demographic characteristics and individual SES variables against the quantiles of the derived SES variable.

We observe that before the ban 21% of the population smoked, with an average of 0.54 packs (or 2.93 individual cigarettes) smoked. These levels of *active* smoking did not fall immediately after the ban. For *passive* smoking, however, we do observe a drop, from a 0.48 probability of being exposed down to 0.26, with the number of hours exposed dropping by 2.26 from 5.28. The average level of cotinine (a biomarker of nicotine levels) also significantly drops. For smoking specifically in pubs, we observe a fall from 7% (34% for smokers) to under 1% (under 3% for smokers). A similar significant change emerges for *passive* smoke exposure in pubs, from 26% of the population to 2%.

In terms of demographic characteristics we observe little difference on average between the samples before and after the ban, with the exception of a small difference for ethnicity. Indeed, by testing if all of these characteristics combined have any explanatory power at explaining if observations are before or after the ban we obtain the F-statistic p-value of 0.211, meaning that one cannot reject the null hypothesis that all β coefficients are equal to zero. Our sample, therefore appears balanced before and after the ban. We nonetheless observe significant differences in the smoking variables.

3.4 Smoking Outcomes Over Time

To illustrate the effect of the ban Figure 2 plots important smoking outcomes across time. The top panel presents *active* smoking variables where the sample is restricted to ever-smokers. The bottom panel shows *passive* smoking outcomes for non-smokers. The scatterplot shows mean outcomes over time, while the white lines and blue areas show the predicted averages and corresponding 95% confidence intervals as estimated by RDD regressions in Tables 3, 4 and 6. To illustrate the long term trend of these smoking outcomes, data from additional waves of the Health Survey for England (waves 2004 until 2010) are used.[16] Within each graph the RDD neighbourhoods (data from the 2007 wave) are shown, enabling

[16] All seven waves are not used in our main analysis for two reasons. First, the regression discontinuity design requires that only observations close to the ban are used. Second, for most waves of the survey it is not possible to identify the year of the interview for all months. The interview period is 15 months (for the main interview) and while the month of interview is available the year of interview is not. This, for example, means that for observations in January, February or March in the 2006 wave the interview could have been held in either 2006 or 2007. Fortunately, in the 2007 wave one variable identifies of the interview date is after the smoking ban, enabling the identification of the correct year. This reveals that there are 37.3% of January interviews in 2008, with 14.8% and 7.4% for February and March respectively. Without this variable a significant amount of observations would have been misclassified. We therefore, conservatively, drop those months in waves other than 2007 in Figure 2.

a visual comparison of the average treatment effects, estimated later, and the longer term trends.

Figure 2: Smoking Outcomes Over Time: 2004-2010

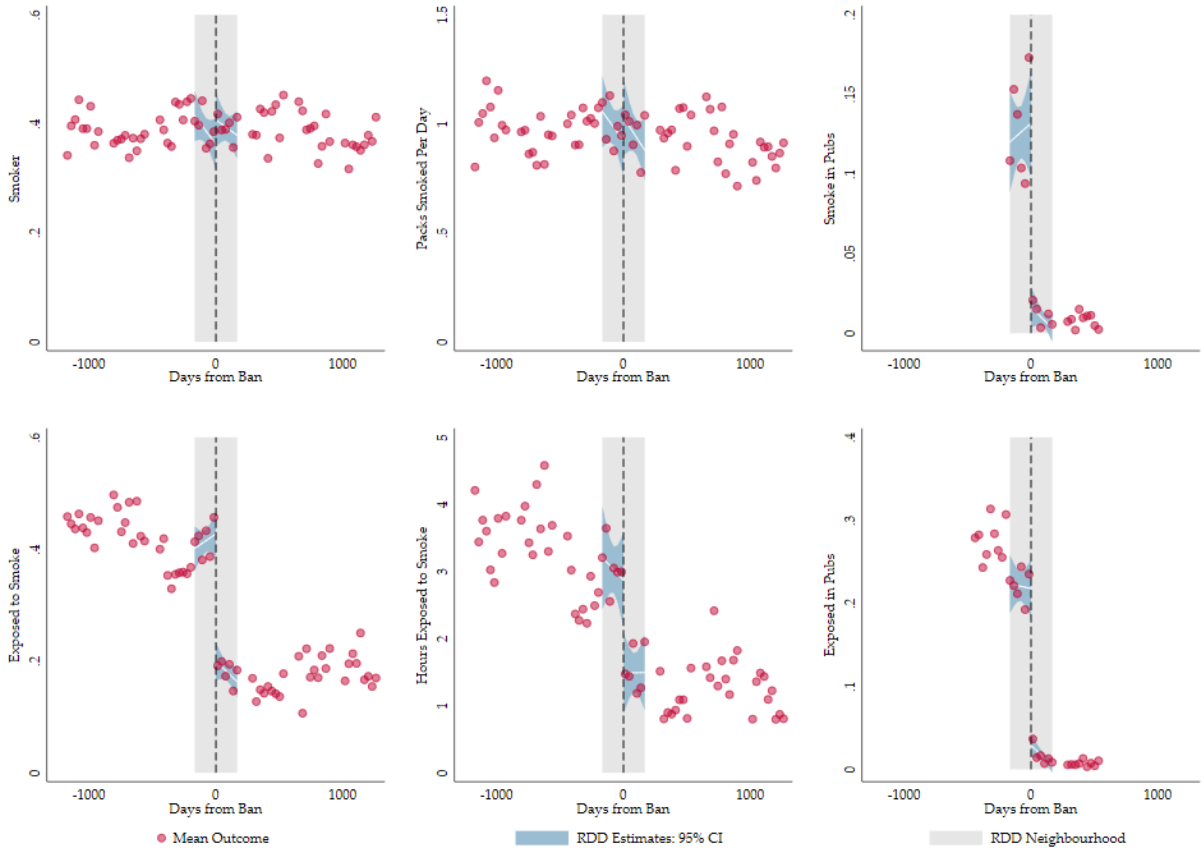

Visually, these plots show that the *level* of active smoking is not reduced by the ban, however there are large, sharp and persistent drops in outcomes relating to passive smoking. Focusing specifically in smoking in pubs, we observe even more stark discontinuities in both active and passive smoking. The exact timing of the ban therefore appears to correspond to large falls in specific smoking outcomes. The following sections focus on the regression discontinuity design, using data within the neighbourhoods shown to estimate the casual treatment effects of the ban.

4 Results

4.1 Average Treatment Effects

Table 3 shows the estimated coefficients from the regression discontinuity design on both *active* (smoker and packs) and *passive* (exposed and hours exposed) smoking. Results show that the smoking ban had no significant short-run effect on the level of active smoking, but caused a large and significant reduction in passive smoking. The average treatment effect for both smoker and packs is not significantly different from zero. However, for non-smokers, the probability of being exposed to second-hand smoke is reduced by 0.228, a 53.4% reduction from 0.427, and the number of hours exposed per week drops from 2.87 to 1.486.

Table 3: Regression Discontinuity Design: Average Treatment Effects

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exposed** Coef./S.E.
Treatment	0.0449	0.1076	-0.2281***	-1.3846***
	(0.0351)	(0.1220)	(0.0266)	(0.4729)
Time	-0.0003	-0.0008	0.0002	-0.0020
	(0.0002)	(0.0009)	(0.0002)	(0.0040)
Treatment X Time	0.0001	-0.0001	-0.0004	0.0021
	(0.0004)	(0.0012)	(0.0003)	(0.0050)
Constant	0.3622***	0.9214***	0.4272***	2.8711***
	(0.0242)	(0.0830)	(0.0203)	(0.3608)
N	3299	3297	4693	4693
R-squared	0.0006	0.0008	0.0660	0.0097
Sample	Ever-Smokers	Ever-Smokers	Non-Smokers	Non-Smokers

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results in Table 3 are robust to a range of sensitivity analyses, which are shown in Appendix A.8. Specific details relating to the bandwidth, functional form and parametric model, alongside placebo timing tests are shown in A.8.1, with three alternative models which include controls in A.8.2. Appendix A.5 additionally analyses cotinine levels, a biomarker of tobacco smoke exposure, and finds results which parallel those above.

4.2 Conditional Average Treatment Effects

While the above shows how the ban affected smoking outcomes on average, it reveals little of the component parts of those average effects. By using our proposed method treatment effects can be estimated conditional on the level of socioeconomic status. Figure 3 plots these

Conditional Average Treatment Effects (CATEs), with corresponding confidence intervals, at each percentile of socioeconomic status for: smoker, packs, exposed and hours exposed.

Figure 3: Conditional Average Treatment Effects

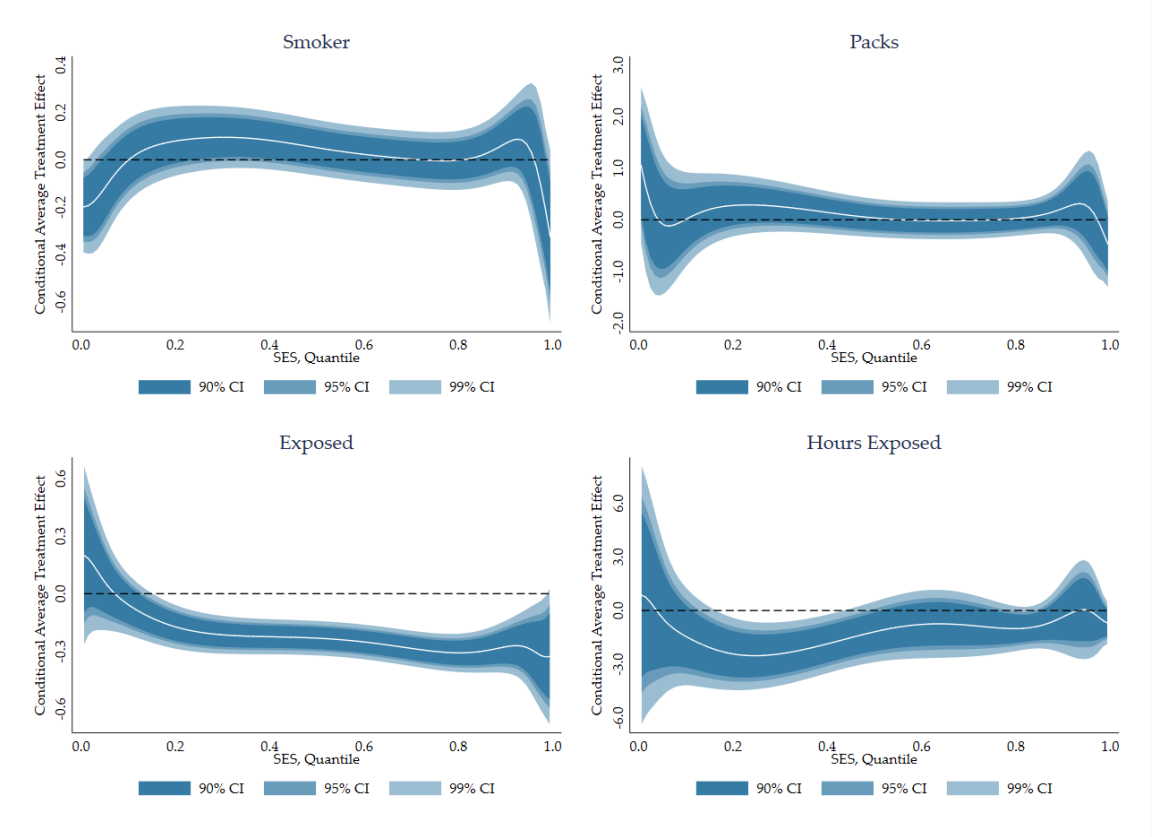

For both smoker and number of packs, results show that there are no significant effects ($p < 0.1$) of the ban at any level of socioeconomic status.[17] For both exposed and hours exposed, however, we observe heterogeneity in the effect size and degree of significance at across the levels of SES. For very low levels of SES there are no significant effects of the ban on either variable. But, above that, we observe opposite trends for exposed and hours exposed. For exposed, the treatment effect increases with SES, while for hours exposed the treatment effect decreases.

Sensitivity analysis relating to the estimated CATEs is provided in Appendix A.9. Alternatives to our socioeconomic status variable are used as the conditional variables (alternative SES, income, IMD, NS-SEC and highest qualification) and different methods (dummy in-

[17]The only exception is that the effect on smoker is negative and significant for those at the extremes. This effect is, however, not robust to alternative RDD specifications (see Appendix A.9.3).

teraction and interaction), weighting functions and precision parameters are explored. The results generally appear to be robust to these alternative specifications.

4.3 Active Smoking

By using more detailed survey data we can identify the location that smokers smoke and whom they smoke near. Using this information we can delve into active smoking, to identify if the ban had an effect on the behaviour of smokers, even if it did not reduce the number of smokers or the amount that they smoke.[18]

Table 4 shows RDD results which reveal the ATEs of the ban on the location where smokers smoke. Results show that the ban did significantly reduce smoking in pubs, with 36.8% of smokers smoking in pubs before the ban reducing to 4.14% immediately after the ban. Reductions were also seen in cafes/shops and in the car. There are no significant reductions at work, but relatively few individuals smoke at work to begin with and the effect size more than halves this. Importantly there is *not* an increase in smoking in the home, or in other's homes. The ban significantly reduced the propensity to smoke in the intended location, the pub, and did not increase the probability that smokers smoked at home.

Table 4: RDD Smoking Location: Smokers

	(1) **Home** Coef./S.E.	(2) **Work** Coef./S.E.	(3) **Other's Home** Coef./S.E.	(4) **Pubs** Coef./S.E.	(5) **Cafes/Shops** Coef./S.E.	(6) **Car** Coef./S.E.
Treatment	-0.0142	-0.0347	-0.0580	-0.3263***	-0.0509**	-0.1005*
	(0.0352)	(0.0302)	(0.0418)	(0.0425)	(0.0238)	(0.0517)
Time	-0.0002	-0.0006**	0.0001	0.0004	-0.0002	0.0001
	(0.0002)	(0.0003)	(0.0003)	(0.0004)	(0.0002)	(0.0004)
Treatment X Time	0.0006*	0.0004	0.0001	-0.0006	0.0001	0.0002
	(0.0003)	(0.0003)	(0.0004)	(0.0004)	(0.0002)	(0.0005)
Constant	0.8854***	0.0724***	0.1717***	0.3677***	0.0652***	0.3286***
	(0.0232)	(0.0269)	(0.0316)	(0.0393)	(0.0219)	(0.0381)
N	1259	1259	1259	1259	1259	1259
R-squared	0.0029	0.0506	0.0028	0.1618	0.0315	0.0059
Sample	Smokers	Smokers	Smokers	Smokers	Smokers	Smokers

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[18]These outcomes are self-reported, hence, what we observe as a reduction in *active* smoking could in fact be an increase in the propensity to misreport the amount of smoking after the ban. The ban could be interpreted as a strong signal that smoking is considered antisocial by the public. This could lead smokers to report lower smoking after the ban. While we cannot address this issue directly, we can use the self-reported outcomes relating to *passive* smoking, as such social stigma should not make non-smokers under-report how much they are exposed after the ban. Section 4.4 shows these results and shows that the ATEs for the location of exposure are generally consistent with the responses from the smokers.

Table 5 reveals the effect of the ban on the propensity of smokers to smoke near 'vulnerable' people. The survey asks if smokers smoke near babies, children, youths, the elderly, pregnant women or those with asthma. Results show that there are significant reductions in the propensity to smoke near all of these groups (bar children). These results imply that the ban had indirect externalities, possibly an educational impact on the harms of second-hand smoke, particularly as some of these groups (babies and youths) are unlikely to be in the pub, the location where smokers have reduced their smoking the most.

Table 5: RDD Smoking Near Whom: Smokers

	(1) **Babies** Coef./S.E.	(2) **Children** Coef./S.E.	(3) **Youths** Coef./S.E.	(4) **Elderly** Coef./S.E.	(5) **Pregnant** Coef./S.E.	(6) **Asthma** Coef./S.E.
Treatment	-0.0434**	-0.0428	-0.0877**	-0.0812**	-0.0461**	-0.0808***
	(0.0183)	(0.0409)	(0.0378)	(0.0405)	(0.0187)	(0.0298)
Time	0.0002	0.0003	0.0004*	0.0000	0.0004***	0.0002
	(0.0001)	(0.0003)	(0.0003)	(0.0003)	(0.0001)	(0.0002)
Treatment X Time	-0.0002	-0.0006	-0.0002	0.0002	-0.0003**	-0.0001
	(0.0002)	(0.0004)	(0.0004)	(0.0004)	(0.0002)	(0.0003)
Constant	0.0541***	0.1845***	0.1686***	0.1753***	0.0539***	0.1161***
	(0.0165)	(0.0303)	(0.0284)	(0.0310)	(0.0161)	(0.0253)
N	1260	1260	1260	1260	1260	1260
R-squared	0.0070	0.0057	0.0050	0.0082	0.0115	0.0116
Sample	Smokers	Smokers	Smokers	Smokers	Smokers	Smokers

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

As shown in Section 4.2 there may be heterogeneity in treatment effects, if estimated conditional on the level of socioeconomic status. This may partially be due to different *levels* of the outcome, conditional on SES, as well as a difference in how effective the treatment has been for individuals with certain levels of SES. One useful output of our proposed analysis is the ability to predict the expected levels of the outcomes, assuming individuals do or do not receive the treatment. This *counterfactual* analysis is particularly useful as it predicts the potential outcomes for the same sample. Within the regression discontinuity design framework this amounts to estimating the full model and predicting outcomes assuming individuals are at the time of the ban ($Z = 0$).

Figure 4 plots these counterfactual predictions, conditional on SES. The outcome of interest is *smoke near vulnerable*, a dummy variable denoting if an individual smokes near any of the six vulnerable groups; with the probability of smoking near vulnerable people being predicted. The model is a standard RDD with control variables. The scatter show the predicted counterfactual outcomes, given the characteristics of that individual belonging to that SES quantile of interest. The line, and 95% confidence intervals, show the expected

counterfactual outcomes for the 'average' individual within the sample.[19] These are further split into subsample of *smokers* and *ever-smokers*, for the left and right panel respectively.

Figure 4: Conditional Counterfactual Predictions: Smoke Near Any Vulnerable Group

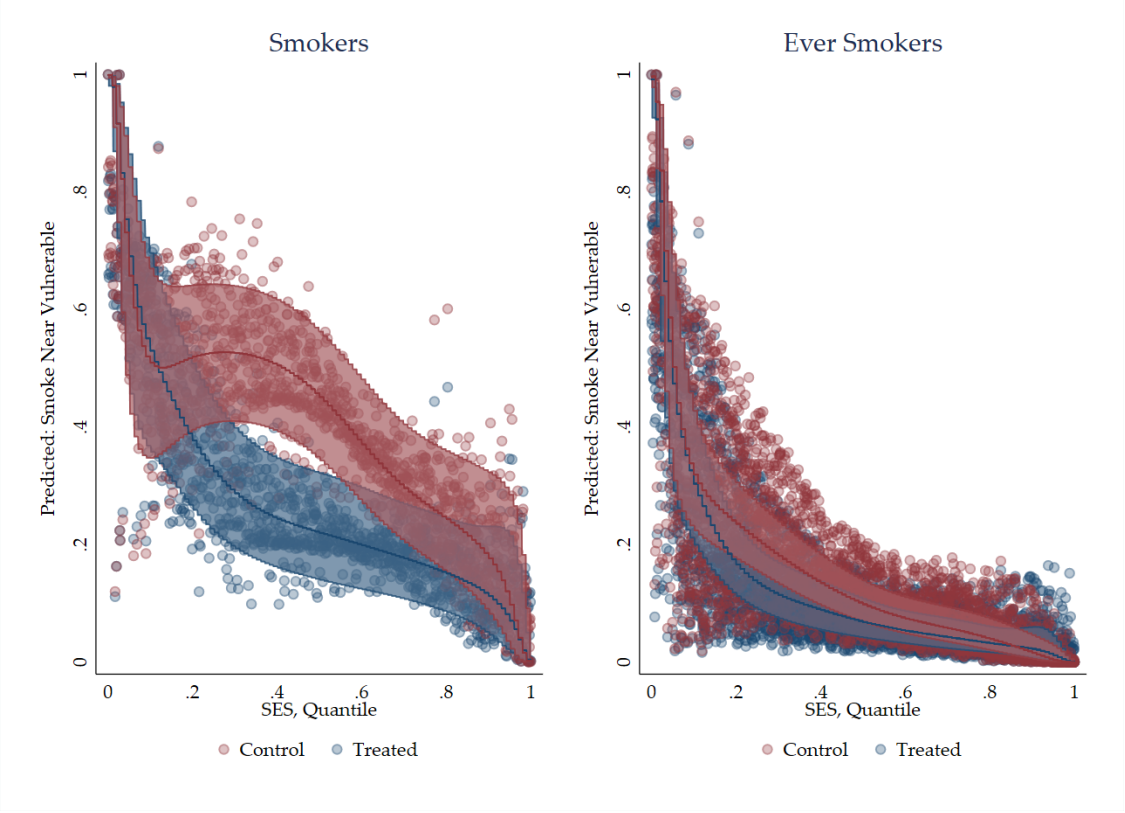

Results show, in both samples, that the predicted probability of smoking near vulnerable groups is higher for those of a low socioeconomic status. The ban was effective in significantly reducing the probability of smoking near vulnerable people for smokers from the 22nd percentile up to the 73rd percentile. However, at the extremes the ban had no significant effect. For those of highest SES the level was already low and there was therefore little potential for reduction. However, for the lowest levels of SES there appears to be no significant reduction, hence the level remains high.

[19]The reasons to show both are: 1) to illustrate the significant differences between control and treatment group, at particular levels of SES; 2) to show how, by including control variables, better predictions about the levels for individuals can be made. While standard RDD models may be suitable for estimating treatment effects, they provide little predictive accuracy. The R^2 for the RDD with controls (for the smokers population) is 0.099 compared to 0.053 of the standard RDD model.

4.4 Passive Smoking

Responses relating to the location of exposure to second-hand tobacco smoke are also available within the survey. Using these responses allows for a more in-depth analysis of the reasons behind the large and significant reductions in passive smoke exposure. Regression discontinuities are run first, to identify ATEs of the ban on the location of exposure, then more sophisticated decomposition analyses are run to identify the contributions of these reductions in reducing the total number of hours exposed.

Table 6 estimates the ATEs of the ban on the location where individuals are exposed to second-hand smoke. We observe large and significant reductions at the pub and no significant change at home. We also observe significant reductions at work, in other's homes, on public transport and in other locations. These results show that the ban significantly reduced exposure to tobacco smoke for non-smokers, particularly in the locations the ban focused upon, and did not have the unintended consequence of increasing exposure at home.

Table 6: RDD Passive Smoke Location: Non-Smokers

	(1) **Home** Coef./S.E.	(2) **Work** Coef./S.E.	(3) **Other's Home** Coef./S.E.	(4) **P. Transport** Coef./S.E.	(5) **Pubs** Coef./S.E.	(6) **Other** Coef./S.E.
Treatment	-0.0137	-0.0313**	-0.0526***	-0.0153***	-0.1892***	-0.0605***
	(0.0134)	(0.0126)	(0.0161)	(0.0057)	(0.0195)	(0.0148)
Time	-0.0001	-0.0000	0.0003**	0.0000	-0.0000	0.0000
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0001)
Treatment X Time	0.0002	-0.0001	-0.0003*	-0.0000	-0.0001	-0.0001
	(0.0001)	(0.0001)	(0.0002)	(0.0001)	(0.0002)	(0.0002)
Constant	0.0483***	0.0612***	0.0984***	0.0169***	0.2179***	0.0892***
	(0.0099)	(0.0104)	(0.0124)	(0.0053)	(0.0181)	(0.0130)
N	4323	4323	4323	4323	4323	4323
R-squared	0.0009	0.0098	0.0049	0.0059	0.1012	0.0224
Sample	Non-Smokers	Non-Smokers	Non-Smokers	Non-Smokers	Non-Smokers	Non-Smokers

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

When comparing the results relating to passive locations to the active smoking locations of Table 4 we observe parallel patterns. Both active and passive variables show a large and significant reduction in pubs and no significant change at home. Responses relating to work and other's home are in the same direction and are of a similar relative magnitude;[20] they are, however, only significant for passive smoke.[21]

[20] A 47.9% reduction in active smoking at work, compared to a 51.1% reduction in passive smoking at work. With a 33.8% reduction at other's homes for active smoking and a 53.5% reduction for passive smoking at other's homes.

[21] This is potentially due to the smaller sample size within the active smokers regressions.

4.4.1 Decomposing Average Treatment Effects

The above separately estimates the effect of the ban on reducing exposure in specific locations. However, as individuals can be exposed in multiple locations, a reduction in one location does not necessarily entail a reduction in general. To assess this, decomposition analysis can be used where the dependant variable is hours exposed and location-specific exposure dummies are used as independent variables. Using this the contribution that a reduction in exposure at a particular location has on a reduction in the total number of hours exposed can be estimated.

Table 7 shows the decomposed ATEs of the ban on the number of hours exposed. The *total* effect is split into the three *component* effects: the endowment, coefficient and unexplained effects. The independent variables are split into the *location exposed* and *control* variables. The interpretation and relevance of the *endowment* and *coefficient* components depends on which type of variables is being considered. Coefficients relating to the interaction of the current variables and time (our forcing variable) are not shown, the coefficients shown can be interpreted as the effects at the time of the ban for the average individual. The total effect, which is equivalent to the ATE, is -1.627 and significant at the 1% level. This effect is almost entirely explained by the *endowment* component, as both the *coefficient* and *unexplained* components are not significant.

The endowment effects of the location exposed variables are the most important here. These results show that, of the total effect of -1.746, the most significant individual components were from a reduction in exposure in pubs (-0.897), at work (-0.491) and at other's homes (-0.137). It is the reduction in individuals being exposed in pubs, at work and at other's homes which explains the majority of the reduction in the total hours exposed. The endowment effects of the control variables are all insignificant and close to zero. In the context of our regression discontinuity design this gives further reassurance that the design is valid. It shows that there is no significant contribution to the total effect, at the time of the ban, from differences in characteristics between control and treatment group.

The coefficient contributions are not significant for any of the location exposed with the exception of being exposed in an other location. The interpretation for this is that the relation that being exposed in an other location has with the number of hours exposed has increased after the ban. No control variables coefficients are significant, meaning that any heterogeneous treatment effects observed explain little of the total (average) effect.

Table 7: Decomposed Average Treatment Effects: Hours Exposed

	(1) Hours Exposed			
	Endowment		***Coefficient***	
	Coef.	S.E.	Coef.	S.E.
Location Exposed				
- Home	-0.121	(0.108)	-0.095	(0.287)
- Work	-0.491***	(0.145)	-0.184	(0.125)
- Other's Home	-0.137*	(0.071)	0.007	(0.163)
- Public Transport	0.022	(0.060)	0.012	(0.011)
- Pubs	-0.897***	(0.216)	-0.021	(0.057)
- Other	-0.122	(0.089)	0.177*	(0.105)
Controls				
Age	0.010	(0.033)	-3.346	(7.728)
Age Squared	0.000	(0.015)	0.958	(3.902)
Sex (1 = female)	0.000	(0.003)	-0.092	(0.465)
Married Couple	0.000	(0.002)	0.135	(0.465)
Ethnicity (1 = white)	-0.001	(0.027)	0.620	(1.146)
Urban	-0.001	(0.003)	1.012	(0.735)
Pregnant	-0.009	(0.017)	0.026	(0.049)
Ever Smoked	0.001	(0.008)	-0.250	(0.382)
Unexplained			1.161	(4.044)
Component Effects	-1.746***	(0.307)	-1.041	(4.023)
Total Effect	-1.627***	(0.440)		
Time Controls	YES			
N	4316			

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.4.2 Decomposing Conditional Average Treatment Effects

The above analysis allows for the decomposition of ATEs into component parts. However, as shown in Section 4.2 treatment effects can vary, conditional on the level of socioeconomic status. By combining these approaches the CATEs can be decomposed into component parts. This allows insights into the potential reasons behind any heterogeneity in CATEs.

Figure 5 shows the decomposed CATEs of the number of hours exposed. The 'total' effect is shown alongside the two components which contribute the most to the total effect. Both are *endowment* effects, showing if individuals are exposed in the pub or at work. From these two components we observe opposite slopes. The exposed in pubs endowment effect, as observed in Table 7, has the largest contribution in reducing the total effect, overall. But we observe an upward slope, in line with the total effect, showing a reduction in the size of effect as SES increases. For exposure at work we observe the opposite trend, its contribution to the total effect is smaller (and indeed insignificant) for the lower levels of SES.

Figure 5: Decomposed CATEs: Hours Exposed (95% CIs)

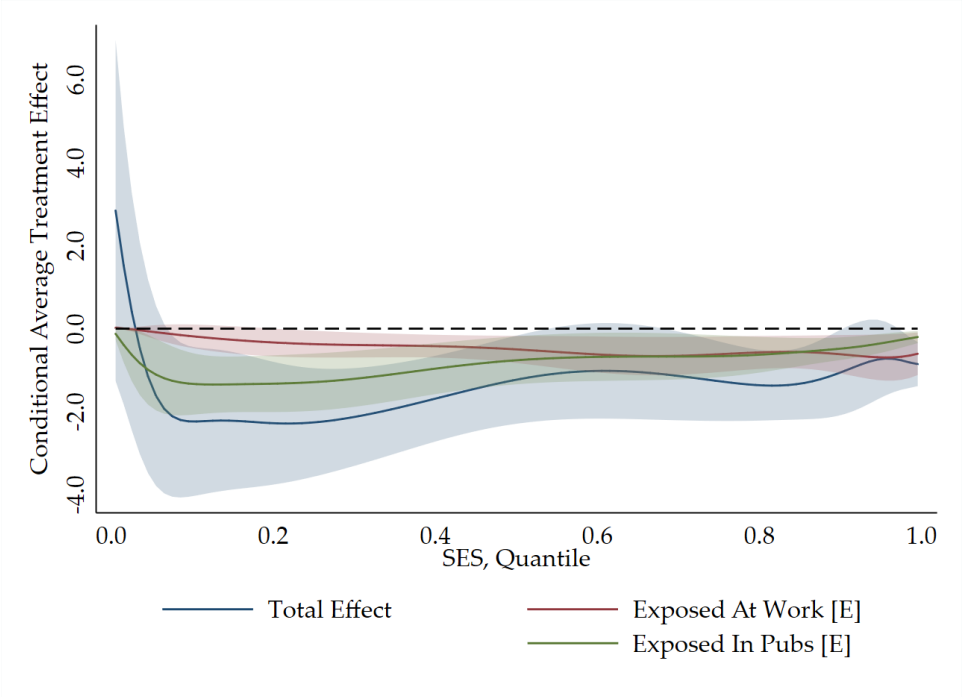

5 Discussion

5.1 Smoking Ban

The smoking ban in England neither reduced smoking prevalence nor intensity of smoking in the short term. The primary aim of the ban, however, was to reduce second-hand exposure to tobacco smoke. In this respect, the smoking ban appears to have been successful: there was a 22.8pp reduction in the prevalence of reported exposure to tobacco smoke for non-smokers following the ban and an average reduction in exposure intensity of 1.38 hours per week. As intended, the ban had its greatest impact in pubs and importantly this was not achieved at the cost of increased exposure in homes.

These results corroborate findings of previous studies, which also reported limited impact of the ban on total tobacco consumption (Fowkes et al., 2008; J. Lee, Glantz, and Millett, 2011; Jones et al., 2015) and substantial reductions in second-hand exposure (Sims et al., 2011). The main additional contribution of our study was to disaggregate the average treatment effects of the smoking ban by socioeconomic status and location, without assuming a particular functional form. For tobacco consumption, we found that the overall lack

of impact of the ban applied across the socioeconomic spectrum. In contrast, the impact of the ban on second-hand exposure varied by both socioeconomic status and location; the greatest effect was in those with the highest baseline exposures towards the more deprived end of the socioeconomic spectrum. For this reason, the legislation appears to have been equity increasing overall, but those with the very lowest socioeconomic status – who also had relatively low baseline exposure – benefited less. In terms of location, there was less of an impact in the workplace for people at the lower end of the socioeconomic spectrum, but a greater impact in pubs.

Our study is subject to several limitations. The Health Survey for England is a large national survey using robust methodology to ensure representativeness, but problems of low numbers emerge once the sample is divided into intersecting groups of interest. This is reflected in the imprecision of estimates for some of our outcomes. We are also reliant on accurate reporting of smoking behaviour and second-hand exposure, which is subject to recall and social desirability bias. Responses may have changed in the periods immediately preceding and following the introduction of the ban, as respondents will have been sensitised by publicity surrounding the legislation, changes to the smoking environment and challenges to their and others' smoking behaviour. These biases could conceivably vary by socioeconomic status, affecting our results. Analysis of cotinine samples, however, corroborates the main results and, although it is possible that respondents to the survey, anticipating cotinine sampling, took additional steps to avoid exposure following the legislation, we think this is unlikely to explain our findings or to limit their generalisability to the general population.

5.2 Methods

The method we develop uses locally-weighted regressions to estimate Conditional Average Treatment Effects (CATEs), when the conditioning variable(s) $\widetilde{\mathbf{x}}$ are continuous. Our method has advantages over the standard methods of dummy interaction (or stratification) and (parametric) interactions, particularly when the relation between CATEs and $\widetilde{\mathbf{x}}$ is complex. As Appendix A.2 shows, when compared to the dummy interaction model our method more accurately estimates CATEs when the number of interactions is small and is more precise when that number is high. When compared to an interaction model, if the degree of parametric interaction assumed is too low our method is much more accurate. The advantage of our method is in its flexibility to allow for any underlying relation between CATEs

and $\widetilde{\mathbf{x}}$; if, however, that relation is shown to be simple then the standard methods may be preferable.

Although the weights used within each local regression can be drawn from the typical kernel functions (Uniform, Normal, Triangular and Epanechnikov) we specifically derive, and use, weights from the Beta distribution. One issue in kernel estimation, with the typical kernel functions, is that coefficient estimates are biased at the bounds. The reason for this is that as the bounds of $\widetilde{x}$ (or $q(\widetilde{x})$) are approached, the weighted average of $\widetilde{\mathbf{x}}$ (or $q(\widetilde{\mathbf{x}})$) moves away from $\widetilde{x}$ (or $q(\widetilde{x})$) in the opposite direction to that bound. The advantage of using the Beta distribution is that this bias does not occur. Precision is still lost at the bounds, but the lack of bias is perhaps advantageous; particularly when the group of interest is at the bounds (i.e. the extreme poor).

One limitation of this paper is the focus on *average* effects. Although *conditional average* treatment effects reveal more than average treatment effect alone, they still seek to estimate the changes in the average of the outcome variable, y, conditional on $\widetilde{x}$. Quantile (Koenker and Bassett Jr, 1978) and distributional regressions (Chernozhukov, Fernández-Val, and Melly, 2013), seek to estimate *distributional* effects, conditional on the quantile (or level) of y, rather than $\widetilde{x}$. Distributional treatment effects show how much y changes at different quantiles of y, providing an analysis which goes 'beyond the mean' (see Fortin, Lemieux, and Firpo (2011) for a review). A further extension to our paper could be to combine these approaches allowing for the estimation of *conditional distributional* treatment effects.

6 Conclusion

Results show that the ban had no immediate impact on the prevalence or intensity of active smoking. However, the ban did reduce passive smoking, through reducing the probability of being exposed and the number of hours of exposure to smoke. Smoking behaviour therefore changed, not in terms of level but in terms of location, with smokers avoiding smoking in particular locations, particularly pubs, and smoking less near vulnerable groups. There was therefore a large reduction in exposure to smoke in pubs, alongside reductions at work, in other's houses and on public transport, with no significant increase in the home.

While these Average Treatment Effects (ATEs) reveal the 'typical' effect of the ban they neglect any heterogeneity in treatment effects and hence ignore any impact the ban has on inequalities in health. To address this we proposed a method for estimating Conditional Av-

erage Treatment Effects (CATEs). Our approach estimates a set of locally-weighted regressions, each centred on quantile(s) of $\widetilde{\mathbf{x}}$, the conditional variable(s) of interest. This approach allows for the estimation of complex relationships between CATEs and $\widetilde{\mathbf{x}}$, decomposition of these CATEs and prediction of counterfactual outcomes. We provide a flexible and easy-to-use Stata program (`lwcate`) to enable others to use our methods for experimental, natural experimental and quasi-experimental designs.

We applied our method to estimate causal effects of the ban, conditional on the level of socioeconomic status (SES). Heterogeneity in CATEs was found for both exposed and hours exposed variables, with greater reductions in hours exposed for those with lower SES. To provide insights into the reasons behind this heterogeneity in CATEs we conducted decomposition analysis. We found that reductions in exposure in the pub and at work contributed most to the overall reduction in the number of hours exposed. But observed that the reduction in hours exposed for the poor was mainly due to a reduction of exposure in the pub, while reductions at both work and the pub contributed to the overall reduction in exposure at higher levels of SES.

These methods to both estimate and decompose Conditional Average Treatment Effects allow for an analysis that goes beyond the mean. The impacts that policies have on inequalities in society can therefore be estimated and better understood, providing more useful and nuanced information for future policy decisions.

References

Abrevaya, Jason, Yu-Chin Hsu, and Robert P Lieli (2015). "Estimating conditional average treatment effects". In: *Journal of Business & Economic Statistics* 33(4), pp. 485–505.

Adda, Jérôme and Francesca Cornaglia (2010). "The effect of bans and taxes on passive smoking". In: *American Economic Journal: Applied Economics* 2(1), pp. 1–32.

Akhtar, Patricia C et al. (2010). "Socioeconomic differences in second-hand smoke exposure among children in Scotland after introduction of the smoke-free legislation". In: *Journal of Epidemiology & Community Health* 64(4), pp. 341–346.

Bedoya, Guadalupe et al. (2017). *Distributional impact analysis: Toolkit and illustrations of impacts beyond the average treatment effect*. The World Bank.

Björklund, Anders and Robert Moffitt (1987). "The estimation of wage gains and welfare gains in self-selection models". In: *The Review of Economics and Statistics*, pp. 42–49.

Blinder, Alan S (1973). "Wage discrimination: reduced form and structural estimates". In: *Journal of Human Resources*, pp. 436–455.

Brown, Tamara, Stephen Platt, and Amanda Amos (2014). "Equity impact of interventions and policies to reduce smoking in youth: systematic review". In: *Tobacco Control* 23(e2), e98–e105.

Calonico, Sebastian et al. (2017). "rdrobust: Software for regression-discontinuity designs". In: *The Stata Journal* 17(2), pp. 372–404.

Chang, Minsu, Sokbae Lee, and Yoon-Jae Whang (2015). "Nonparametric tests of conditional treatment effects with an application to single-sex schooling on academic achievements". In: *The Econometrics Journal* 18(3), pp. 307–346.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly (2013). "Inference on counterfactual distributions". In: *Econometrica* 81(6), pp. 2205–2268.

Cleveland, William S (1979). "Robust locally weighted regression and smoothing scatterplots". In: *Journal of the American Statistical Association* 74(368), pp. 829–836.

Cleveland, William S and Susan J Devlin (1988). "Locally weighted regression: an approach to regression analysis by local fitting". In: *Journal of the American Statistical Association* 83(403), pp. 596–610.

Cornelissen, Thomas et al. (2016). "From LATE to MTE: Alternative methods for the evaluation of policy interventions". In: *Labour Economics* 41, pp. 47–60.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo (2011). *Decomposition Methods in Economics, Volume 4 of Handbook of Labor Economics, Chapter 1*.

Fowkes, Freya JI et al. (2008). "Scottish smoke-free legislation and trends in smoking cessation". In: *Addiction* 103(11), pp. 1888–1895.

Hahn, Jinyong (1998). "On the role of the propensity score in efficient semiparametric estimation of average treatment effects". In: *Econometrica*, pp. 315–331.

Heckman, James J and Richard Robb Jr (1985). "Alternative methods for evaluating the impact of interventions: An overview". In: *Journal of Econometrics* 30(1-2), pp. 239–267.

Heckman, James, Hidehiko Ichimura, and Petra Todd (1998). "Matching as an econometric evaluation estimator". In: *The Review of Economic Studies* 65(2), pp. 261–294.

Heckman, James, Hidehiko Ichimura, and Petra E Todd (1997). "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme". In: *The Review of Economic Studies* 64(4), pp. 605–654.

Hsu, Yu-Chin (2017). "Consistent tests for conditional treatment effects". In: *The Econometrics Journal* 20(1), pp. 1–22.

Imbens, Guido W and Joshua D Angrist (1994). "Identification and Estimation of Local Average Treatment Effects". In: *Econometrica* 62(2), pp. 467–475.

Imbens, Guido W and Thomas Lemieux (2008). "Regression discontinuity designs: A guide to practice". In: *Journal of econometrics* 142(2), pp. 615–635.

Jann, Ben (2008). "The Blinder–Oaxaca decomposition for linear regression models". In: *The Stata Journal* 8(4), pp. 453–479.

Jarvis, Martin J et al. (2012). "Impact of smoke-free legislation on children's exposure to secondhand smoke: cotinine data from the Health Survey for England". In: *Tobacco Control* 21(1), pp. 18–23.

Jones, Andrew et al. (2015). "Do public smoking bans have an impact on active smoking? Evidence from the UK". In: *Health Economics* 24(2), pp. 175–192.

Khan, Shakeeb and Elie Tamer (2010). "Irregular identification, support conditions, and inverse weight estimation". In: *Econometrica* 78(6), pp. 2021–2042.

Koenker, Roger and Gilbert Bassett Jr (1978). "Regression quantiles". In: *Econometrica: Journal of the Econometric Society*, pp. 33–50.

Lee, John, Stanton A Glantz, and Christopher Millett (2011). "Effect of smoke-free legislation on adult smoking behaviour in England in the 18 months following implementation". In: *PloS One* 6(6), e20933.

Lee, Sokbae, Ryo Okui, and Yoon-Jae Whang (2017). "Doubly robust uniform confidence band for the conditional average treatment effect function". In: *Journal of Applied Econometrics* 32(7), pp. 1207–1225.

Machado, José AF and José Mata (2005). "Counterfactual decomposition of changes in wage distributions using quantile regression". In: *Journal of Applied Econometrics* 20(4), pp. 445–465.

MaCurdy, Thomas, Xiaohong Chen, and Han Hong (2011). "Flexible estimation of treatment effect parameters". In: *American Economic Review* 101(3), pp. 544–51.

Melly, Blaise (2005). "Decomposition of differences in distribution using quantile regression". In: *Labour Economics* 12(4), pp. 577–590.

Moore, Graham, Dorothy Currie, et al. (2012). "Socioeconomic inequalities in childhood exposure to secondhand smoke before and after smoke-free legislation in three UK countries". In: *Journal of Public Health* 34(4), pp. 599–608.

Moore, Graham, Jo Holliday, and Laurence Moore (2011). "Socioeconomic patterning in changes in child exposure to secondhand smoke after implementation of smoke-free legislation in Wales". In: *Nicotine & Tobacco Research* 13(10), pp. 903–910.

Nadaraya, Elizbar A (1964). "On estimating regression". In: *Theory of Probability & Its Applications* 9(1), pp. 141–142.

Oaxaca, Ronald (1973). "Male-female wage differentials in urban labor markets". In: *International Economic Review*, pp. 693–709.

Pagan, Adrian and Aman Ullah (1999). *Nonparametric econometrics*. Cambridge university press.

Rubin, Donald B (1974). "Estimating causal effects of treatments in randomized and nonrandomized studies." In: *Journal of Educational Psychology* 66(5), p. 688.

Sims, Michelle et al. (2011). "Did smokefree legislation in England reduce exposure to secondhand smoke among nonsmoking adults? Cotinine analysis from the Health Survey for England". In: *Environmental health perspectives* 120(3), pp. 425–430.

Słoczyński, Tymon (2015). "The Oaxaca–Blinder unexplained component as a treatment effects estimator". In: *Oxford Bulletin of Economics and Statistics* 77(4), pp. 588–604.

Thistlethwaite, Donald L and Donald T Campbell (1960). "Regression-discontinuity analysis: An alternative to the ex post facto experiment." In: *Journal of Educational Psychology* 51(6), p. 309.

Vyas, Seema and Lilani Kumaranayake (2006). "Constructing socio-economic status indices: how to use principal components analysis". In: *Health Policy and Planning* 21(6), pp. 459–468.

Watson, Geoffrey S (1964). "Smooth regression analysis". In: *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 359–372.

A Appendix

A.1 Kernel, Local-Linear and Local-Polynomial Estimation

While the main method sets out the estimation procedure for a non-parametric estimation which closely resembles a kernel regression approach, it is straight forward to extend this model to a local-linear or local-polynomial estimation. The main model estimates M locally-weighted regressions at j quantiles of $\widetilde{\mathbf{X}}$, $q_j^*(\widetilde{\mathbf{X}})$:

$$\sqrt{\mathbf{w}_j}\,\mathbf{y} = \sqrt{\mathbf{w}_j}\left(\boldsymbol{\beta}_{0j} + \mathbf{X}\boldsymbol{\beta}_j + \mathbf{u}\right) \tag{18}$$

This method non-parametrically estimates coefficients $\hat{\boldsymbol{\beta}}_j$, for all individuals i who belong to quantile $q_j^*(\widetilde{\mathbf{X}})$. Meaning that all individuals within that quantile have the same "constant" coefficient estimates, as in a kernel regression. This restriction can be lifted by estimating the set of M local-linear regressions:

$$\sqrt{\mathbf{w}_j}\,\mathbf{y} = \sqrt{\mathbf{w}_j}\left(\boldsymbol{\beta}_{0j} + \mathbf{X}\boldsymbol{\beta}_j + (\widetilde{\mathbf{X}}_j - \widetilde{\overline{\mathbf{X}}}_j)\boldsymbol{\beta}_{Lj} + \mathbf{u}\right) \tag{19}$$

The difference between the kernel and local-linear estimation is that, for the kernel estimation $\widetilde{\mathbf{X}}$ is not explicitly included within the model while in the local-linear estimation the difference between $\widetilde{\mathbf{X}}_j$ and the mean of $\widetilde{\mathbf{X}}_j$ within a particular quantile j, $\widetilde{\overline{\mathbf{X}}}_j$, is included. This allows not only for "constant" coefficients to be estimated, for all individuals within quantile j, but a linear relation for individual i from their distance to the mean $\widetilde{\mathbf{X}}_j$ of their quantile, $\widetilde{\overline{\mathbf{X}}}_j$. This can be further extended to a local-polynomial regression, where the parametric form of this line is not assumed to be linear, where D is the degree of the polynomial chosen:

$$\sqrt{\mathbf{w}_j}\,\mathbf{y} = \sqrt{\mathbf{w}_j}\left(\boldsymbol{\beta}_{0j} + \mathbf{X}\boldsymbol{\beta}_j + \sum_{d}^{D}(\widetilde{\mathbf{X}}_j - \widetilde{\overline{\mathbf{X}}}_j)^d\boldsymbol{\beta}_{Ldj} + \mathbf{u}\right) \tag{20}$$

These methods are particularly useful for the typical weighting functions within the non-parametric literature. As at bounds the weighted average rank (or X) generated using each of the Uniform, Triangular, Normal and Epanechnikov distributions is biased away from that bound. This, however, is less of an issue for the weights derived from the Beta distribution.

A.2 Simulation Testing

In order to test the proposed method we simulate data for 200,000 individuals. We model health as being related to socioeconomic status (SES), time and treatment status. We model five alternative functional forms, where we vary the relationship between the treatment effects and SES. The underlying parameters are known, enabling the calculation of analytical conditional average treatment effects (CATEs), for each individual i within the sample.

We test how well our model performs by calculating *error*, the mean squared-error between the individual-level estimated and analytical CATEs, and *precision*, the standard errors of the estimates. Our method will be compared to alternative methods: dummy interaction models and (parametric) interaction models. Then the model choices within our method will be tested: the precision parameter, number of quantiles, kernel weighting function and type of local estimator. Finally, Monte-Carlo simulations will be run to establish how well the model performs with finite samples. The structure of the simulated data is:

$$y_i = \beta_0 + \widetilde{x}_i\beta_1 + f_\tau(\widetilde{x}_i, \beta_\tau)d_i + z_i\beta_2 + (z_i * d_i)\beta_3 + u_i \tag{21}$$

Where, for individual i, y_i denotes health, $\widetilde{x}_i$ socioeconomic status (SES), the treatment dummy is d_i, time is denoted by z_i and the error is u_i. SES is randomly drawn from the Beta distribution, $\widetilde{x}_i \sim Beta(3, 2.5)$, while the error term is normally distributed, $u \sim N(0, 1) * 2$. Time z_i is uniformly drawn, between -100 and 0, for those in the control and between 0 and 100 for the treated. The constant $\beta_0 = 50$, while $\beta_1 = 10$, meaning that health is positively associated with SES. For simplicity $\beta_2 = \beta_3 = 0$, to ensure constant time trends.

Our main interest is in the function $f_\tau(\widetilde{x}_i, \beta_\tau)$, which denotes the conditional average treatment effect. This function is different for each of the five alternative simulations we run and it is this we aim to estimate and compare. The five functional forms are as follows:

$$A \qquad No\,Relation : f_T(.) = 5$$

$$B \qquad Linear : f_T(.) = 10.\widetilde{x}_i$$

$$C \qquad Quadratic : f_T(.) = 25.\widetilde{x}_i - 25.\widetilde{x}_i^2$$

$$D \qquad Cubic : f_T(.) = -15.\widetilde{x}_i + 95.\widetilde{x}_i^2 - 85.\widetilde{x}_i^3$$

$$E \qquad Discontinuous : f_T(.) = 2.\widetilde{x}_i + 5.I[\widetilde{x}_i > 0.5] + 5.I[\widetilde{x}_i > 0.75]$$

The resulting simulations for 200,000 people are shown in Figure 6.

Figure 6: Simulated Data: Alternative Functional Forms

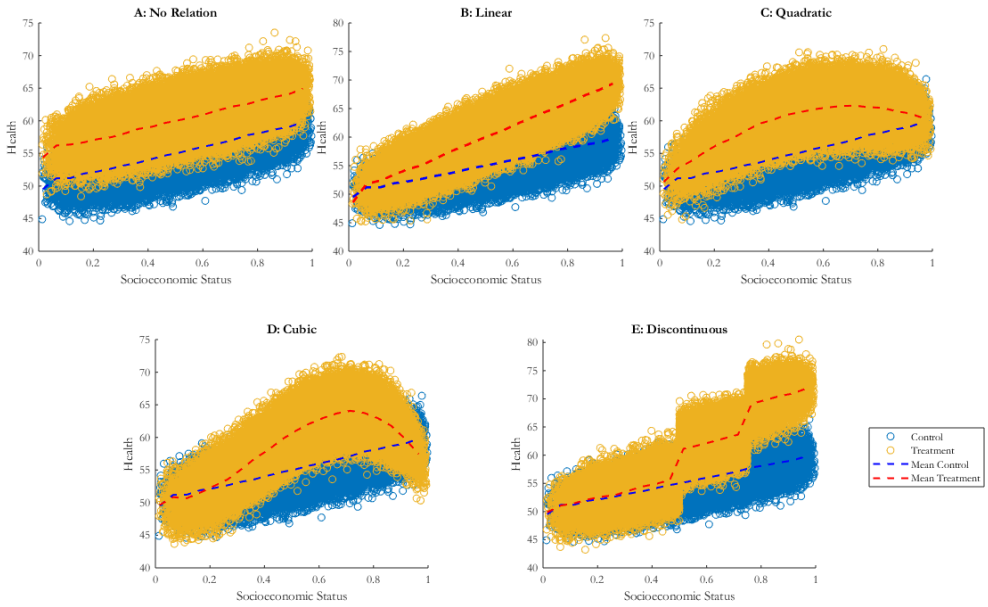

Table 8 shows the error (Mean Squared Error) and precision (Standard Error) of the Beta method, in comparison to dummy interaction models, where the number of interacted dummies of $\widetilde{x}_i$ is varied. Error in estimation generally increases as the complexity of the 'true' model increases. For the dummy interaction models with low numbers of interactions this increase is rapid, while the increase for the 100 interactions model is kept low. The error for the Beta model is comparable to the 100 interactions model, with the exception of 'true' Model E. Precision is constant across the 'true' models, but decreases as the number of interactions increase. The Beta method consistently lies between the 5 and 10 dummy interaction models. This simulation shows that the error of the Beta model is comparable to a dummy interaction with 100 interactions (with the exception of Model E), but it has precision similar to models with far fewer interactions.

Table 9 compares the Beta method to (parametric) interaction models, where the degree of the interaction with $\widetilde{x}_i$ changes. The degree of interaction (Constant, Linear, Quadratic and Cubic) correspond directly to the 'true' models (A. B, C and D). This is apparent when observing the low error on the upper diagonal. If the degree of the interaction is general enough, then the interaction models do well and outperform the Beta method. However, when the degree of interaction is not general or the 'true' model is not parametric (Model

E) enough then interaction models have much higher levels of error than the Beta method does. Precision decreases as the degree of interaction increases, but is consistently lower than the Beta method. Table 9 shows that if the relation between CATEs and $\widetilde{x}_i$ is simple the interaction models are preferable, but if they are too complicated the Beta method is preferred. As this relation will always be unknown a priori, the Beta method may be useful to test the complexity of the relation.

Table 8: Dummy Interaction Comparison

	Mean Squared Error					
	2	5	10	20	100	Beta
A	0.001	0.003	0.005	0.009	0.033	0.011
B	1.196	0.259	0.083	0.033	0.035	0.023
C	1.175	0.425	0.170	0.068	0.035	0.032
D	3.684	1.282	0.503	0.191	0.054	0.058
E	5.330	1.011	0.810	0.486	0.101	0.727
	Standard Error					
	2	5	10	20	100	Beta
A	0.030	0.044	0.061	0.085	0.190	0.055
B	0.035	0.045	0.061	0.085	0.190	0.056
C	0.031	0.045	0.061	0.086	0.191	0.055
D	0.035	0.046	0.062	0.086	0.190	0.057
E	0.041	0.047	0.063	0.087	0.191	0.058

Table 9: Interaction Comparison

	Mean Squared Error				
	Constant	Linear	Quadratic	Cubic	Beta
A	0.001	0.001	0.002	0.003	0.011
B	3.807	0.001	0.001	0.002	0.023
C	1.283	1.178	0.001	0.002	0.032
D	5.174	3.638	0.741	0.003	0.058
E	15.586	2.475	1.930	1.715	0.727
	Standard Error				
	Constant	Linear	Quadratic	Cubic	Beta
A	0.019	0.026	0.031	0.035	0.055
B	0.021	0.026	0.031	0.035	0.056
C	0.020	0.028	0.031	0.035	0.055
D	0.023	0.031	0.032	0.035	0.057
E	0.027	0.029	0.034	0.038	0.058

Tables 10 and 11 illustrate how the error and precision of our method changes when the precision parameter, s, and the number of quantiles, M, are varied. Table 10 shows that as the s increases error is reduced, across all 'true' models. Precision, on the other hand, has a U-shaped trend. Precision begins to improve as s increases, but after a point precision starts to worsen. Table 11 shows that when there is no parametric 'true' relation between CATE and $\widetilde{x}_i$ M has little effect, but when the parametric relation is complex then higher M becomes important. Precision generally seems to be little affected by M. As a high number of M is computationally intensive a lower M is desirable, but this is perhaps advisable only when the relation between $\widetilde{x}_i$ M is simple.

The above show simulation results when our method used weights derived from the Beta Distribution and uses a kernel estimator. Alternative weights and estimators can be used (and are available in our Stata code). Table 12 compares estimations which use the Beta Distribution, to more standard kernel weighting functions, Uniform, Triangular, Normal and Epanechnikov. Results show that the Beta generally is less prone to error, but has lower

Table 10: Alternative Precision, s

	Mean Squared Error				
	5	10	50	100	500
A	0.109	0.052	0.011	0.007	0.009
B	0.302	0.133	0.023	0.011	0.009
C	0.760	0.264	0.032	0.016	0.012
D	1.737	0.570	0.058	0.031	0.026
E	1.890	1.563	0.727	0.511	0.240
	Standard Error				
	5	10	50	100	500
A	0.099	0.067	0.055	0.058	0.081
B	0.115	0.073	0.056	0.058	0.081
C	0.110	0.071	0.055	0.058	0.081
D	0.118	0.075	0.057	0.059	0.081
E	0.107	0.071	0.058	0.060	0.083

Table 11: Different Quantiles, M

	Mean Squared Error				
	10	20	50	100	500
A	0.004	0.004	0.006	0.011	0.016
B	0.082	0.029	0.016	0.023	0.028
C	0.171	0.071	0.033	0.032	0.032
D	0.510	0.205	0.082	0.058	0.052
E	1.006	0.816	0.740	0.727	0.730
	Standard Error				
	10	20	50	100	500
A	0.045	0.047	0.051	0.055	0.055
B	0.046	0.048	0.052	0.056	0.056
C	0.046	0.047	0.052	0.055	0.055
D	0.046	0.048	0.053	0.057	0.056
E	0.049	0.050	0.054	0.058	0.057

levels of precision.[22] Table 13 shows alternative forms of local estimator. Results show that both error and precision are relatively constant across these different estimation models.[23]

Table 12: Alternative Kernel Weights

	Mean Squared Error				
	Beta	Uniform	Triang.	Normal	Epan.
A	0.011	0.002	0.003	0.002	0.002
B	0.023	0.231	0.089	0.052	0.119
C	0.032	0.488	0.244	0.181	0.308
D	0.058	1.441	0.700	0.548	0.888
E	0.727	2.575	1.358	1.494	1.659
	Standard Error				
	Beta	Uniform	Triang.	Normal	Epan.
A	0.055	0.032	0.034	0.033	0.033
B	0.056	0.035	0.036	0.035	0.035
C	0.055	0.033	0.035	0.034	0.034
D	0.057	0.036	0.037	0.036	0.036
E	0.058	0.041	0.040	0.039	0.040

Table 13: Local Estimator

	Mean Squared Error				
	Kernel	Linear	LinearInt	Quadratic	QuadInt
A	0.011	0.010	0.013	0.009	0.027
B	0.023	0.020	0.016	0.019	0.030
C	0.032	0.029	0.021	0.029	0.034
D	0.058	0.056	0.038	0.055	0.047
E	0.727	0.733	0.738	0.731	0.498
	Standard Error				
	Kernel	Linear	LinearInt	Quadratic	QuadInt
A	0.055	0.054	0.054	0.054	0.060
B	0.056	0.054	0.054	0.054	0.059
C	0.055	0.054	0.054	0.054	0.059
D	0.057	0.055	0.054	0.055	0.059
E	0.058	0.056	0.055	0.056	0.061

Finally, Table 14 uses a Monte-Carlo simulation to assess how our method performs with small samples. The results, generally, are as expected the higher the sample size the lower the error and the more precise the estimates are.

[22] As different measure of precision are used between these methods more analysis needs to be undertaken to assess if this trend holds more generally.

[23] Perhaps where more complex estimation models would be useful is where lower M is desirable.

Table 14: Sample Size

	Mean Squared Error					
	50	100	500	1000	5000	10000
A	6.519	2.903	0.598	0.342	0.113	0.069
B	6.775	3.058	0.625	0.347	0.106	0.076
C	6.774	3.071	0.646	0.361	0.137	0.106
D	7.368	3.274	0.694	0.427	0.177	0.144
E	8.053	3.984	1.385	1.059	0.822	0.778
	Standard Error					
	50	100	500	1000	5000	10000
A	1.683	1.308	0.654	0.475	0.222	0.162
B	1.724	1.335	0.663	0.482	0.225	0.164
C	1.702	1.319	0.659	0.479	0.225	0.163
D	1.747	1.348	0.671	0.488	0.228	0.167
E	1.847	1.417	0.698	0.506	0.235	0.169

A.3 Alternative Beta Formulations

A.3.1 Variance

Within our method to estimate CATEs we make a particular assumption in relation to variance. Here, we show an alternative assumption which could also be used. Rather than assuming variance is dependent upon q_j^*, we assume constant variance, as follows:

$$E[Q_j] = \frac{a_j}{a_j + b_j} = q_j^* \tag{22}$$

$$Var[Q_j] = \frac{a_j b_j}{(a_j + b_j)^2(a_j + b_j + 1)} = s \tag{23}$$

It follows that:

$$a_j = q_j^* \left(q_j^* \left(\frac{1 - q_j^*}{s} \right) - 1 \right), \qquad b_j = \left(q_j^* \left(\frac{1 - q_j^*}{s} \right) - 1 \right) (1 - q_j^*) \tag{24}$$

This formulation ensures a constant variance; which could be a desirable property, particularly approaching the bounds. However, it runs into the issue that for particular values of q_j^* and s the values of a_j and b_j are negative. A constraint of the Beta distribution is that $a_j, b_j > 0$, meaning that this formulation is often not feasible. An issue which does not emerge with the assumption in the main model. If, however, this formulation is used it is sufficient to ensure: $s < q_j^*(1 - q_j^*)$. This can be done by either selecting lower values of s (the variance) or putting bounds on the q_j^* at which the weights are centred on.

A.3.2 Dirichlet

While the Beta distribution specified focuses on one dimension at a time, the Dirichlet distribution is multivariate beta distribution, which could be useful for determining weights when the number of dimensions, k, is greater than one. In the paper and Stata code we multiply different Beta distributions to establish weights, however, this is primarily due to computational limitations because the CDF of the Dirichlet is not easily estimated in Stata. It could, however, prove useful for future research. The formulation is as follows:

$$E[Q_k] = \frac{a_k}{a_0} = q_k^* \tag{25}$$

$$Var(Q_k) = \frac{(a_k(a_0 - a_k))}{(a_0^2(a_0 + 1))} = \frac{(q_k^*(q_0^* - q_k^*))}{s} \tag{26}$$

Where:

$$a_0 = \sum_{k=1}^{K} a_k, \qquad q_0^* = \sum_{k=1}^{K} q_k^*$$

It follows that, $\forall k$:

$$q_k^*(s-1) = a_k \tag{27}$$

The resulting weights would be as follows:

$$w_{ij} = \int_{V_j} \cdots \int \left(\frac{1}{B(a_{0j})} \prod_{k=1}^{K_j} \ddot{q}_{kj}^{a_{kj}-1} \right) d\ddot{q}_{1j} \ldots d\ddot{q}_{\kappa j} \tag{28}$$

Where:

$$B(a_{0j}) = \frac{\prod_{k=1}^{K_j} \Gamma(a_{kj})}{\Gamma\left(\sum_{k=1}^{K_j} a_{kj}\right)}, \qquad \ddot{q}_{K_j} = 1 - \sum_{k=1}^{\kappa_j} \ddot{q}_{kj},$$

$$V_j = \left\{ (\ddot{q}_{1j}, \ldots, \ddot{q}_{\kappa j}) \in \mathbf{R}^{\kappa_j} \ : \ r_{ik} - \frac{1}{2N} \leq \ddot{q}_{kj} \leq r_{ik} + \frac{1}{2N}, \forall k \in [1, \kappa_j] \right\}$$

These weights are calculated for each individual, accounting for K dimensions of variables (i.e. SES, age and weight), where $\kappa = K - 1$, at each j_k quantile we wish to estimate the weighted least-squares model at.

A.4 Theory: RDD Decomposition

To show intuitively how decomposition analysis can be used with an RDD approach we start with the basics. The standard RDD model is as follows:

$$\mathbf{y} = \beta_0 + \mathbf{D}\beta_1 + \mathbf{Z}\beta_2 + (\mathbf{D}' \times \mathbf{Z})\beta_3 + \mathbf{u} \tag{29}$$

From this we can establish the Average Treatment Effect (ATE) at the discontinuity, when $\mathbf{Z} \to 0$. For $E[\mathbf{y}|\mathbf{Z} = 0]$, the equation is:

$$E[\mathbf{y}|\mathbf{Z} = 0] = \beta_0 + E[\mathbf{D}]\beta_1 \tag{30}$$

β_0 is the expected value of y for the control group, when $\mathbf{Z} = 0$, while β_1 is the expected difference between control and treatment group (the ATE), when $\mathbf{Z} = 0$. The expected y at the discontinuity is then: $\beta_0 + \beta_1 E[\mathbf{D}]$; the expected value of y for the control group, plus the ATE multiplied by the proportion of those who are treated.

This same analysis can be conducted using a decomposition approach. With two separate equations, where the subscripts denote the control (0) and treated (1) group:

$$\mathbf{y}_0 = \gamma_{00} + \mathbf{Z}_0\gamma_{01} + \mathbf{u}_0 \tag{31}$$

$$\mathbf{y}_1 = \gamma_{10} + \mathbf{Z}_1\gamma_{11} + \mathbf{u}_1 \tag{32}$$

Comparing the two methods, $\gamma_{00} = \beta_0$, $\gamma_{10} = \beta_0 + \beta_1$, $\gamma_{01} = \beta_2$ and $\gamma_{11} = \beta_2 + \beta_3$. To obtain the ATE the difference between control and treatment group, at the discontinuity, needs to be taken:

$$E[\mathbf{y_1}|\mathbf{Z}_{0,1} = 0] - E[\mathbf{y_0}|\mathbf{Z}_{0,1} = 0] = \gamma_{10} - \gamma_{00} = \beta_1 \tag{33}$$

The benefit of the decomposition method is that we can decompose the 'total' ATE into component parts if we include X variables and their interactions with Z:

$$\mathbf{y}_0 = \delta_{00} + \mathbf{Z}_0\delta_{01} + \mathbf{X}_0\delta_{0j} + (\mathbf{Z}_0' \times \mathbf{X}_0)\delta_{0k} + \mathbf{u}_0 \tag{34}$$

$$\mathbf{y}_1 = \delta_{10} + \mathbf{Z}_1\delta_{11} + \mathbf{X}_1\delta_{1j} + (\mathbf{Z}_1' \times \mathbf{X}_1)\delta_{1k} + \mathbf{u}_1 \tag{35}$$

It follows that the ATE can be obtained as above. Then, if this equation is rearranged (by subtracting and adding $E[\mathbf{X}_1]\delta_{0j}$ from the third and fourth term, respectively) the distinct *coefficient*, *endowment* and *unexplained* components can be observed:

$$\begin{aligned} E[\mathbf{y_1}|\mathbf{Z}=0] - E[\mathbf{y_0}|\mathbf{Z}=0] &= \delta_{10} - \delta_{00} + E[\mathbf{X}_1]\delta_{1j} - E[\mathbf{X}_0]\delta_{0j} \\ &= \underbrace{\delta_{10} - \delta_{00}}_{Unexplained} + \underbrace{E[\mathbf{X}_1](\delta_{1j} - \delta_{0j})}_{Coefficients} + \underbrace{(E[\mathbf{X}_1] - E[\mathbf{X}_0])\delta_{0j}}_{Endowments} \end{aligned} \tag{36}$$

The ATE (at the discontinuity) is thus decomposed into three components.[24] This is for the average, in the sense that it looks to the average individual, but it concerns the treatment effect at the time of the ban.

A.4.1 Two-Stage Extension

The above allows the average treatment effect to be decomposed. However, one issue emerges which relates to the particular X variables, the 'mechanism' variables X_M. By trying to identify potential mechanisms through which the smoking ban reduced the total number of hours exposed we include dummies indicating exposure in particular locations (i.e at home and in the pub). Within the decomposed model the *endowment* component identifies this effect. However, it identifies the difference between the expected values of the $X's$ ($E[\mathbf{X}_{M1}] - E[\mathbf{X}_{M0}]$) rather than the difference in the expected values, conditional on $\mathbf{Z}=0$ ($E[\mathbf{X}_{M1}|\mathbf{Z}=0] - E[\mathbf{X}_{M0}|\mathbf{Z}=0]$). The latter would allow for a causal interpretation of the contribution to the ATE that the 'mechanism' variables have. In other words, that the difference in expectation of the X_M variables is caused by the intervention and this causal difference is what contributes to the ATE.

In order to achieve this we propose a two-stage procedure. 1) Estimate $\hat{\mathbf{X}}_M|Z=0$; 2) Plug in $\hat{\mathbf{X}}_M|Z=0$ and $\mathbf{X}_M - \hat{\mathbf{X}}_M|Z=0$ into the decomposition equation, in place of $\mathbf{X}_M$. As follows:

[24] One issue with interpretation of the decomposition is that both the *unexplained* and *coefficient* components are determined by how $\mathbf{X}$ is centred. The interpretation of the *unexplained* component has to be the unexplained component at the point where all $\mathbf{X}=0$. If, then, $\mathbf{X}$ is mean centred about zero, for whatever reason, then both the *unexplained* and *coefficient* components will change; the sum of those two components and the *endowment* component will, however, remain the same. In our analysis (particularly as the sign of the *coefficient* can easily switch if the $\mathbf{X}$ are mean centred) we do not mean centre, leaving the interpretation as follows. The *unexplained* component is where all $\mathbf{X}=0$, while the *coefficient* components can be interpreted at their mean. The exception to this is the forcing variable Z, which has to be centred on zero; meaning that the *unexplained* component related to the average treatment effect at the time of the ban and not elsewhere.

Stage 1:

Separately estimate regressions for each 'mechanism' variable $\mathbf{X}_M$. With explanatory variables D, the treatment dummy, Z, time, $\mathbf{X}_C$, the control variables and their interactions:

$$\mathbf{X}_M = \beta_0 + \mathbf{D}\beta_1 + \mathbf{Z}\beta_2 + (\mathbf{D}' \times \mathbf{Z})\beta_3 + \mathbf{X}_C\beta_4 + \mathbf{u} \tag{37}$$

Predict $\hat{\mathbf{X}}_M|\mathbf{Z} = 0$, for each 'mechanism' variable.

Stage 2:

Run the decomposition analysis, plugging in the predicted variables, $\hat{\mathbf{X}}_M|\mathbf{Z} = 0$, and the difference between the predicted values at $Z = 0$ and the observed 'mechanism' variables, $\mathbf{X}_M - \hat{\mathbf{X}}_M|Z = 0$, to estimate the control and treatment equations separately:[25]

$$\begin{aligned}\mathbf{y}_0 = \lambda_{00} + \mathbf{Z}_0\lambda_{01} + \mathbf{X}_{C0}\lambda_{0j} + (\mathbf{Z}_0' \times \mathbf{X}_{C0})\lambda_{0k} + (\hat{\mathbf{X}}_{M0}|\mathbf{Z}_0 = 0)\lambda_{0m} + \\ (\mathbf{X}_{M0} - \hat{\mathbf{X}}_{M0}|\mathbf{Z}_0 = 0)\lambda_{0n} + (\mathbf{Z}_0' \times (\mathbf{X}_{M0} - \hat{\mathbf{X}}_{M0}|\mathbf{Z}_0 = 0))\lambda_{0p} + \mathbf{u}_0\end{aligned} \tag{38}$$

$$\begin{aligned}\mathbf{y}_1 = \lambda_{10} + \mathbf{Z}_1\lambda_{11} + \mathbf{X}_{C1}\lambda_{1j} + (\mathbf{Z}_1' \times \mathbf{X}_{C1})\lambda_{1k} + (\hat{\mathbf{X}}_{M1}|\mathbf{Z}_1 = 0)\lambda_{1m} + \\ (\mathbf{X}_{M1} - \hat{\mathbf{X}}_{M1}|\mathbf{Z}_1 = 0)\lambda_{1n} + (\mathbf{Z}_1' \times (\mathbf{X}_{M1} - \hat{\mathbf{X}}_{M1}|\mathbf{Z}_1 = 0))\lambda_{1p} + \mathbf{u}_1\end{aligned} \tag{39}$$

Then decompose the expected differences between the control and treatment group, at the discontinuity. Where, for notational ease: $ATE = E[\mathbf{y_1}|\mathbf{Z} = 0] - E[\mathbf{y_0}|\mathbf{Z} = 0]$, $(\hat{\mathbf{X}}_M|\mathbf{Z} = 0) = \widehat{\mathbf{X}}_M$ and $(\mathbf{X}_{M0} - \hat{\mathbf{X}}_M|\mathbf{Z} = 0) = \widetilde{\mathbf{X}}_M$:

$$\begin{aligned}ATE &= \lambda_{10} - \lambda_{00} + E[\mathbf{X}_{C1}]\lambda_{1j} - E[\mathbf{X}_{C0}]\lambda_{0j} + \\ &\quad E[\widehat{\mathbf{X}}_{M1}]\lambda_{1m} - E[\widehat{\mathbf{X}}_{M0}]\lambda_{0m} + E[\widetilde{\mathbf{X}}_{M1}]\lambda_{1n} - E[\widetilde{\mathbf{X}}_{M0}]\lambda_{0n} \\ &= \underbrace{\lambda_{10} - \lambda_{00}}_{Unexplained} + \underbrace{E[\mathbf{X}_{C1}](\lambda_{1j} - \lambda_{0j})}_{C.Coefficient} + \underbrace{(E[\mathbf{X}_{C1}] - E[\mathbf{X}_{C0}])\lambda_{0j}}_{C.Endowment} + \\ &\qquad \underbrace{E[\widehat{\mathbf{X}}_{M0}](\lambda_{1m} - \lambda_{0m})}_{M.Coefficient} + \underbrace{(E[\widehat{\mathbf{X}}_{M1}] - E[\widehat{\mathbf{X}}_{M0}])\lambda_{0m}}_{M.Endowment} + \\ &\qquad \underbrace{E[\widetilde{\mathbf{X}}_{M1}](\lambda_{1n} - \lambda_{0n})}_{D.Coefficient} + \underbrace{(E[\widetilde{\mathbf{X}}_{M1}] - E[\widetilde{\mathbf{X}}_{M0}])\lambda_{0n}}_{D.Endowment}\end{aligned} \tag{40}$$

[25]Both are included to ensure that the linear combination of the two $\hat{\mathbf{X}}_M|\mathbf{Z} = 0 + \mathbf{X}_M - \hat{\mathbf{X}}_M|Z = 0 = \mathbf{X}_M$. This allows data not to be lost and the effect that difference in the predicted 'mechanism' variable at the discontinuity (interpreted as being caused by the smoking ban) has been separated out.

For the purpose of estimating the *endowment* component, the *mechanism endowment* is the component of the ATE which can be explained by the difference between the 'mechanism' variables, where this difference was caused by the intervention.

While this two-stage extension could help predict the *causal mechanism* components of the ATE, it is reliant upon good estimation of $\widehat{\mathbf{X}}_M$. In our case, our predictions of $\widehat{\mathbf{X}}_M$ are not precise; meaning that this approach is perhaps not suitable. Moreover, as shown in Table 6, time (the forcing variable) has very little effect on the 'mechanism' variables, meaning that $E[\mathbf{X}_M] \approx E[\hat{\mathbf{X}}_M|\mathbf{Z} = 0]$, negating the need for the two step procedure.

A.5 Biomarker: Cotinine

The outcomes used in the main analysis are self-reported measures from the Health Survey for England. One issue when individuals self-report is the potential for misreporting; individuals may over or under report the outcome meaning that the observed outcome is different from the true outcome. If this misreporting is not consistent between individuals and across time, our estimated treatment effects could simply be due to a shift in misreporting, rather than a genuine change in the true value.

Conveniently, within the survey a more objective measure of tobacco smoke exposure is available. Cotinine levels are measured in a follow up nurses visit, derived from saliva samples of respondents. Cotinine is the primary metabolite of nicotine and is a biomarker of tobacco smoke exposure. By using cotinine, rather than self-reported levels of tobacco smoke, we can identify if the observed treatment effects persist. One issue with the level of continine, however, is that we cannot distinguish between active and passive smoke exposure for those who are smokers (nor can we distinguish between cotinine derived from tobacoo exposure or other sources of nicotine, such as replacement therapy). If we identify a drop in cotinine levels, due to the ban, for smokers, we cannot identify if it is due to a drop in active or passive smoke exposure. We can, however, identify this for non-smokers.

Table 15 shows results from the regression discontinuity design, which identifies the effect of the smoking ban on cotinine levels for different populations: smokers, ever-smokers, ex-smokers, non-smokers and all. The results are comparable to the treatment effects for the self-reported outcomes, in Table 3. We observe that there is a significant drop in second-hand smoke exposure, both for hours exposed and cotinine levels, for non-smokers. As mentioned above, for smokers we cannot precisely disentangle active and passive smoke exposure; but the results within different subgroups are suggestive of similar results in Table 3. In Table 3, we observe no effect of the ban on active smoking, here we similarly observe no significant effect of the ban on cotinine levels for smokers. Similar results emerge for ever-smokers, ex-smokers and all; where the effects are always negative, but not significant. This is suggestive of a drop in passive smoke, but little reduction in cotinine levels through active smoking.

There are several issues with the cotinine variable which means we do not use it in our main analysis. Firstly, the nurses visit only took place for a subset of the sample (approximately 27.5% of the adult sample did not complete the nurses survey) and of those interviewed 83.6% had a valid cotinine measure. Secondly, the nurses visits usually took place after the main interview. While most took place within a month, the self-reported questions

Table 15: RDD: Effects of the Ban on Cotinine Levels

	(1) **Cotinine** Coef./S.E.	(2) **Cotinine** Coef./S.E.	(3) **Cotinine** Coef./S.E.	(4) **Cotinine** Coef./S.E.	(5) **Cotinine** Coef./S.E.
Treatment	-33.9534	-17.4802	-9.5398	-5.6673**	-10.8572
	(26.4602)	(15.8461)	(9.2572)	(2.5189)	(9.2656)
Time	0.0043	-0.1336	-0.1208	0.0156	-0.1124
	(0.1995)	(0.1221)	(0.0739)	(0.0225)	(0.0739)
Treatment X Time	0.0658	0.2524	0.2231**	0.0335	0.2153**
	(0.2667)	(0.1643)	(0.0989)	(0.0279)	(0.0988)
Constant	307.5716***	114.0129***	59.8176***	6.0500***	61.0719***
	(17.4393)	(10.4726)	(6.1453)	(2.2843)	(6.1678)
N	774	2066	3736	2975	3749
R-squared	0.0054	0.0029	0.0022	0.0022	0.0022
Sample	Smokers	Ever-Smokers	Ex-Smokers	Non-Smokers	All

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

in the main interview are usually time dependent, meaning outcomes could potentially have changed. The final issue relates to the extreme values of the variable. Cotinine levels follow a poisson like distribution, with a modal spike of cotinine is at 0 (30.56%), but with a maximum value of 986.2. These extreme values could bias the OLS estimates. To address this, here, we run both *poisson* and *negative binomial distribution* regressions, and run the RDD using $log(cotinine + 1)$; we find the same direction and significance levels of our results as in Table 15, for each sub-population.

A.6 Missing Income

Of 6,331 adults 1,421 observations are missing for equivalised income (22.45%) in 2007. To recover these observations, we impute income for those missing. To do so, we regress log equivalised income (henceforth income) on: household size, number of children, age, age squared, sex, marriage status, ethnicity, urban, IMD quintile, savings/investments, benefits, NS-SEC, Labour Market status, highest qualification, housing tenure and overcrowding. We then predict income out-of-sample for the missing observations.

The left panel of Figure 7 shows the distribution of both observed ($N = 4,910$) and imputed ($N = 1,421$) income. The graph highlights that the distributions are similar, but as it typical with single imputation methods, the variance is lower for the imputed data. While this could be an issue if income alone were used in our analysis, we use income as only one of 31 variables to estimate socioeconomic status, using the principal component analysis. As shown in the right panel of Figure 7 the distribution of SES for imputed and not as very similar, due to the additional variation of the other 30 variables. Indeed, weights derived for SES are very similar when using or not using the observations with imputed income. As a result, we use the 1,421 individuals within the main analysis to increase sample size.

Figure 7: Log Equivalised Income and SES: Missing Income

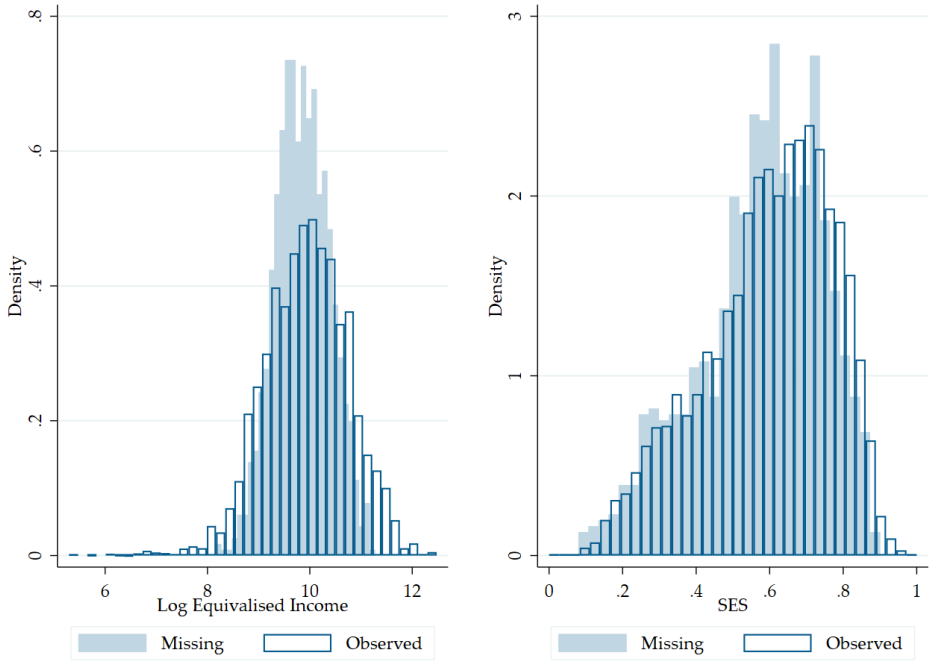

To show robustness to the inclusion of imputed data, the main analysis for the ATEs and CATEs are shown below. Table 16 shows that the significance levels of the ATEs are the same and the magnitudes are not significantly different from the main analysis. Figure 8 shows the the CATEs estimates are also very similar. Hence, in order to maintain statistical power for our analyses we use the imputed data where necessary.

Table 16: RDD: Average Treatment Effects, No Imputed Income

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exposed** Coef./S.E.
Treatment	0.0406	0.1404	-0.2610***	-1.8355***
	(0.0393)	(0.1377)	(0.0296)	(0.5157)
Time	-0.0004	-0.0010	0.0002	-0.0017
	(0.0003)	(0.0010)	(0.0002)	(0.0042)
Treatment X Time	0.0005	0.0001	-0.0001	0.0054
	(0.0004)	(0.0014)	(0.0003)	(0.0056)
Constant	0.3570***	0.9169***	0.4412***	2.9864***
	(0.0265)	(0.0912)	(0.0224)	(0.3872)
N	2711	2709	3792	3792
R-squared	0.0010	0.0008	0.0711	0.0106
Sample	Ever-Smokers	Ever-Smokers	Non-Smokers	Non-Smokers

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 8: Conditional Average Treatment Effects, No Imputed Income

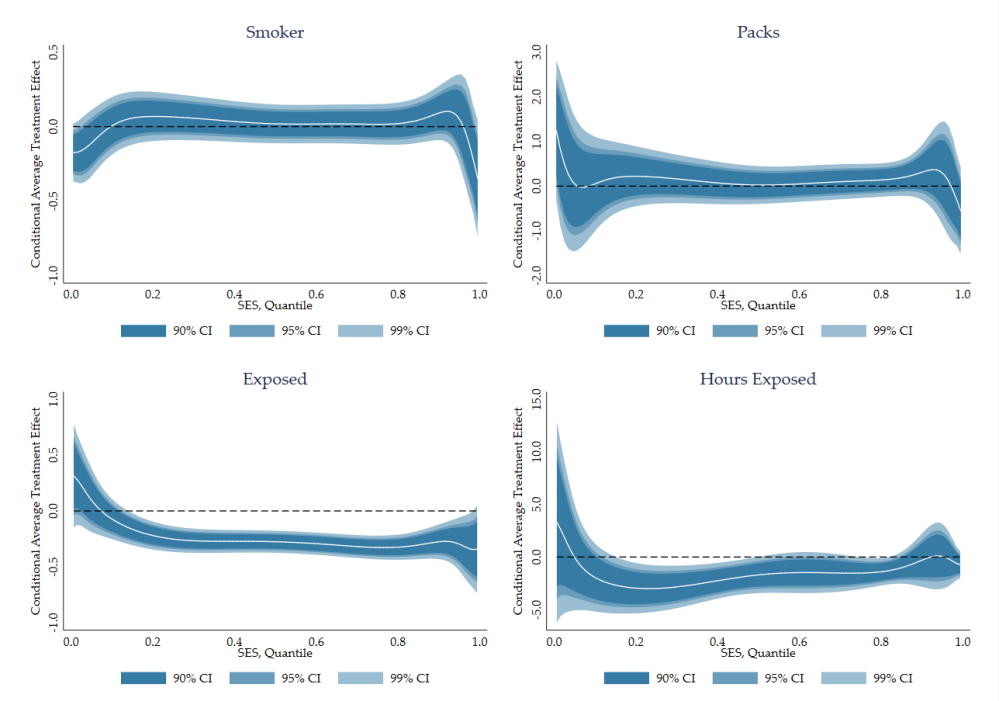

A.7 Socioeconomic Status: Descriptive Statistics

Table 17 provides descriptive statistics which relate to our derived index of socioeconomic status (SES). It summarises the mean values of demographic characteristics and individual SES variables across the whole sample, conditional on a quantile of SES. Table 18 show this information for the smoking outcomes for the control group only, split into smokers, ever-smokers and non-smokers. These two tables highlight the relation between smoking outcomes and SES, alongside the sample composition of different quantiles of SES. For demographic characteristics this highlights who is in each quantile, while for the SES variables it reveals how these variables interact to determine SES.

For demographics, we observe interesting trends, which indicate who has higher and lower SES. As the measure is estimated conditional on age, the age variable is expected to be relatively constant, however, we still observe particularly looking at the extremes 45 compared to 52, that those who are older, within their age group tend to have higher SES. Females make up 67% in the poorest 5% compared to 42% in the highest 5%. Similarly those who are poor are more likely to be single, non-white, from an urban setting and pregnant.

For the SES variables, while the SES is constructed with them, it is of interest to see how our measure relates to these important variables and how they interact with one another. To give the examples of the extremes. Those in the lowest 5% have the lowest income and live in the areas with the lowest IMD. They receive no income from savings, but 90% receive benefits. No individuals are in the managerial/professional class, while 38% are in routine occupations and 24% have never been employed, with only 1% currently being employed. There are 72% with no qualifications and none with a degree, while 96% of the group rent. This is in stark contrast to the highest 5%, who earn the most and live in the best areas. Those with earnings from savings are 68%, with less than 1% receiving any form of benefits. 91% are in managerial/professional occupations, with 99% either currently employed or retired. Most, 72% have a degree, while 99% own their house.

The strength of this measure lies in the identification of those characteristics which are important in determining SES. Moreover, it allows for the interactions and amalgamations of these variables which determine where each individual is placed on a continuous index.

Table 18 shows, in general, there is a strong socioeconomic gradient for smoking. Those with lower SES are more likely to smoke, be exposed and have been smokers in the past, they also smoke more and are exposed for more hours. These self-reported measures follow

Table 17: Socioeconomic Status: Descriptive Statistics

	Low		**Quintiles**				**High**		**Summary**
	≤ 0.05 Mean	≤ 0.1 Mean	**Q1** Mean	**Q2** Mean	**Q3** Mean	**Q4** Mean	≥ 0.9 Mean	≥ 0.95 Mean	**Mean** Mean
Demographics									
Age	45.38	49.26	49.84	52.29	48.99	49.57	50.61	52.28	50.35
Sex (1 = female)	0.67	0.66	0.65	0.59	0.52	0.48	0.44	0.42	0.56
SES	0.20	0.25	0.35	0.55	0.67	0.80	0.84	0.86	0.59
Married Couple	0.33	0.36	0.46	0.61	0.62	0.64	0.67	0.69	0.58
Ethnicity (1 = white)	0.81	0.83	0.84	0.86	0.92	0.95	0.95	0.96	0.90
Urban	0.92	0.91	0.88	0.80	0.74	0.69	0.68	0.66	0.77
Pregnant	0.03	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01
SES Variables									
Log Equivalised Income	8.93	9.03	9.25	9.74	10.21	10.74	10.88	10.99	9.98
IMD Quintile	1.65	1.80	2.11	2.92	3.51	3.94	4.24	4.46	3.13
Savings & Investments	0.00	0.00	0.02	0.10	0.21	0.40	0.56	0.68	0.17
No Benefits	0.10	0.18	0.36	0.73	0.81	0.97	0.99	0.99	0.72
NS-SEC									
- Higher Manager/Profess.	0.00	0.00	0.01	0.02	0.08	0.29	0.40	0.45	0.10
- Lower Manager/Profess.	0.00	0.01	0.04	0.13	0.30	0.50	0.49	0.46	0.23
- Intermediate	0.03	0.06	0.07	0.16	0.20	0.10	0.06	0.06	0.13
- Small Employ./Own Acc.	0.03	0.03	0.06	0.11	0.12	0.05	0.02	0.02	0.09
- Lower Superv./Tech.	0.05	0.08	0.10	0.13	0.12	0.01	0.01	0.00	0.08
- Semi-routine Occup.	0.25	0.26	0.31	0.26	0.12	0.03	0.01	0.01	0.19
- Routine Occup.	0.38	0.37	0.29	0.16	0.04	0.01	0.01	0.01	0.13
- Never Employed	0.24	0.17	0.11	0.02	0.01	0.00	0.00	0.00	0.03
Labour Market Status									
- Employed	0.01	0.05	0.22	0.53	0.69	0.75	0.73	0.71	0.55
- Unemployed	0.06	0.05	0.04	0.01	0.00	0.00	0.00	0.00	0.01
- Retired	0.13	0.25	0.27	0.32	0.23	0.22	0.25	0.28	0.27
- Full Time Student	0.01	0.02	0.05	0.05	0.04	0.01	0.00	0.00	0.03
- Looking After Home	0.42	0.33	0.23	0.06	0.03	0.02	0.01	0.01	0.09
- Other Inactive	0.37	0.29	0.18	0.02	0.01	0.00	0.00	0.00	0.05
Highest Qualification									
- Degree	0.00	0.01	0.03	0.09	0.18	0.50	0.66	0.72	0.20
- Higher Education	0.02	0.02	0.04	0.08	0.14	0.16	0.15	0.13	0.11
- A Level	0.02	0.03	0.07	0.09	0.20	0.13	0.08	0.05	0.12
- GCSE A*-C	0.16	0.15	0.19	0.26	0.27	0.18	0.10	0.09	0.22
- GCSE D-G	0.07	0.06	0.08	0.05	0.05	0.01	0.00	0.00	0.05
- Other	0.01	0.01	0.01	0.02	0.01	0.01	0.00	0.00	0.02
- None	0.72	0.72	0.58	0.41	0.15	0.02	0.00	0.00	0.29
Housing Tenure									
- Own Outright	0.03	0.04	0.14	0.41	0.37	0.38	0.44	0.48	0.34
- Own - Mortgage/Loan	0.00	0.03	0.14	0.35	0.51	0.60	0.55	0.51	0.39
- Part Rent/Part Mort.	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
- Rent	0.96	0.92	0.71	0.22	0.09	0.02	0.00	0.00	0.25
- Live Rent Free	0.00	0.00	0.01	0.02	0.03	0.00	0.01	0.01	0.01
Overcrowding	1.25	1.19	1.17	1.01	0.90	0.73	0.67	0.63	0.94

the same trend as the cotinine. This negative relation is monotonic, with the exception that the highest 5% seem to have higher levels of smoking outcomes, than the top 10%.

By moving between the panels of Table 18 we disentangle the social gradient of active and passive smoking. Conditional on being a smoker, we observe U-shaped patterns. The poorest quartile have higher smoking outcomes than the second quartile, but the third quartile

have lower smoking outcomes than the richest. Indeed, the richest 5% have the highest levels of cotinine. For ever-smnokers the steep monotone trend re-emerges when looking at quartiles, but remains U-shaped when considering the richest 5%. Conditional on being a non-smokers, we observe no relation between SES and the probability of being exposed or having ever smoked. However, the number of hours exposed and cotinine levels maintain a socioeconomic gradient. Conditional on being a smoker, the rich are much more likely to smoke and be exposed in pubs; for non-smokers, however, the trend is less clear.

Table 18: Socioeconomic Status: Smoking Descriptive Statistics for Control Group

	Low		**Quintiles**				**High**		**Summary**
	≤ 0.05 Mean	≤ 0.1 Mean	**Q1** Mean	**Q2** Mean	**Q3** Mean	**Q4** Mean	≥ 0.9 Mean	≥ 0.95 Mean	**Mean** Mean
All									
Smoker	0.59	0.45	0.38	0.20	0.16	0.10	0.07	0.06	0.21
Packs Smoked Per Day	1.80	1.37	1.11	0.49	0.33	0.21	0.14	0.17	0.54
Exposed to Smoke	0.62	0.54	0.53	0.45	0.48	0.44	0.42	0.39	0.48
Hours Exposed to Smoke	11.85	9.02	8.71	4.64	4.22	2.70	2.08	2.18	5.09
Ever Smoked	0.77	0.73	0.70	0.59	0.58	0.54	0.54	0.53	0.60
Cotinine Levels	222.29	167.96	138.01	72.20	51.84	26.73	24.81	27.05	69.79
Smoke in Pubs	0.13	0.10	0.11	0.05	0.06	0.05	0.04	0.04	0.07
Exposed in Pubs	0.23	0.18	0.22	0.23	0.28	0.28	0.25	0.22	0.25
Smokers									
Smoker	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Packs Smoked Per Day	3.04	3.04	2.89	2.43	2.07	2.15	2.04	2.67	2.55
Exposed to Smoke	0.76	0.71	0.73	0.64	0.69	0.79	0.92	1.00	0.71
Hours Exposed to Smoke	16.95	14.96	15.70	9.77	10.12	11.74	11.13	16.58	12.77
Ever Smoked	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Cotinine Levels	329.84	337.36	342.91	312.71	277.56	279.35	353.07	438.50	314.81
Smoke in Pubs	0.22	0.23	0.29	0.25	0.39	0.54	0.61	0.58	0.33
Exposed in Pubs	0.30	0.26	0.31	0.28	0.41	0.59	0.65	0.50	0.35
Ever-Smokers									
Smoker	0.77	0.62	0.55	0.34	0.28	0.18	0.13	0.12	0.35
Packs Smoked Per Day	2.34	1.89	1.59	0.83	0.57	0.39	0.26	0.32	0.90
Exposed to Smoke	0.71	0.62	0.61	0.49	0.55	0.49	0.47	0.45	0.54
Hours Exposed to Smoke	14.52	11.34	11.18	5.65	5.66	3.38	2.51	3.02	6.78
Ever Smoked	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Cotinine Levels	273.39	221.58	191.80	113.58	84.99	49.43	45.27	48.53	112.79
Smoke in Pubs	0.17	0.14	0.16	0.08	0.10	0.10	0.08	0.07	0.11
Exposed in Pubs	0.30	0.23	0.28	0.24	0.33	0.34	0.29	0.26	0.29
Non-Smokers									
Smoker	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Packs Smoked Per Day	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Exposed to Smoke	0.42	0.40	0.41	0.40	0.44	0.40	0.39	0.35	0.41
Hours Exposed to Smoke	4.61	4.24	4.43	3.33	3.10	1.73	1.40	1.20	3.04
Ever Smoked	0.44	0.50	0.51	0.49	0.50	0.49	0.50	0.50	0.50
Cotinine Levels	25.10	11.18	13.64	10.52	10.11	3.74	4.48	1.93	8.92
Smoke in Pubs	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Exposed in Pubs	0.15	0.11	0.17	0.21	0.25	0.24	0.21	0.20	0.22

A.8 Sensitivity of the Regression Discontinuity Design

This appendix shows sensitivity analyses which evaluate the robustness of the baseline results of Table 3. Sensitivity analyses relating to bandwidth, functional form, model specification and the inclusion of controls are conducted, alongside placebo timing tests.

A.8.1 Bandwidths, Functional Form, Model and Placebo Timing

While the regression discontinuity design is generally considered to have high levels of internal validity (i.e. a high level of confidence that the observed effects are indeed causal) sensitivity analysis should be conducted to identify if: A) the bandwidth chosen affects the estimated effects, B) the functional form assumed changes the estimated effects, C) results are sensitive to the regression model used and D) if similar effects exist if the discontinuity is moved. This section will address these concerns for the main four variables.

Figure 9: Sensitivity of Average Treatment Effects: Forest Plots

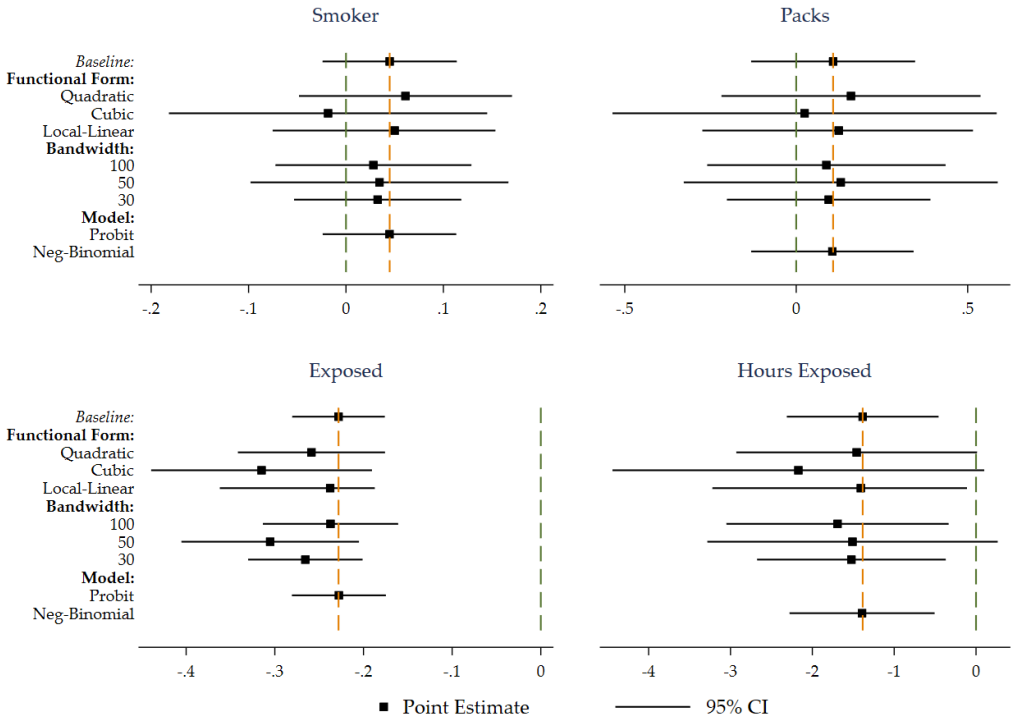

Figure 9 shows forest plots of ATEs for: smokers, packs, exposed and hours exposed. Each plot shows results from each of the sensitivity analyses. Across all four variables we observe there are no significant differences between the baseline and any of the sensitivity analyses.

For both smoker and packs we always observe that estimates are not significantly different from zero. While effects for both exposed and hours exposed are consistently negative and significant at the 5% level (with the exception of quadratic and bandwidth: 50 which are still significant at the 10% level). These results show the direction and significance of the baseline ATEs are robust to a host of sensitivity analyses.

The first sensitivity analysis involves moving the bandwidth within which the RDD is estimated. In the main specification the bandwidth of 190 days (January to December of 2007) used. However, if the RDD is valid the estimated effect should remain stable for alternative bandwidth. Table 19 shows the ATEs for each of the main variables, at bandwidths of 100, 50 and 30 days. Across all three bandwidths the sign, significance levels and magnitudes remain similar across all four variables.

Table 19: Bandwidth Sensitivity

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exp** Coef./S.E.
TE (BW = 100)	0.0282	0.0880	-0.2372***	-1.6923**
	(0.0513)	(0.1772)	(0.0389)	(0.6923)
N	1737	1736	2485	2485
TE (BW = 50)	0.0344	0.1298	-0.3052***	-1.5090*
	(0.0675)	(0.2333)	(0.0510)	(0.9045)
N	1126	1125	1641	1641
TE (BW = 30)	0.0325	0.0946	-0.2655***	-1.5225***
	(0.0437)	(0.1511)	(0.0329)	(0.5872)
N	506	505	739	739

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

One issue with the RDD is that ATEs are often sensitive to the functional form assumed. While our main model assumes a linear trend through time, it is possible that this trend is non-linear. Assuming the trend is non-linear could, therefore, result in a biased estimate of the ATE. To test this the functional form across time is relaxed. Table 20 show results from three models. The first assume the trend over time is quadratic, while the second assumes it is cubic. The third estimates the trend using a local-linear approach, a non-parametric method (from the *rdrobust* command (Calonico et al., 2017)) which makes no assumptions about the functional form of the time trend. Across each model, we obtain similar results to our base model. The largest change appears to be for the cubic model, for hours exposed, however, this effect is not significantly different from the other estimates.

Table 20: Functional Form Sensitivity

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exp** Coef./S.E.
- Quadratic TE	0.0611	0.1597	-0.2587***	-1.4583*
	(0.0558)	(0.1927)	(0.0423)	(0.7502)
N	3299	3297	4693	4693
- Cubic TE	-0.0184	0.0244	-0.3149***	-2.1707*
	(0.0833)	(0.2857)	(0.0635)	(1.1584)
N	3299	3297	4693	4693
- Local Linear TE	0.0500	0.1241	-0.2375***	-1.4074***
	(0.0376)	(0.1302)	(0.0286)	(0.5003)
N	3299	3297	4693	4693

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Ordinary least squares (OLS) is used throughout the main analysis. The reason for this is for ease of interpretation of coefficients. However, other regression models may be better suited at fitting the underlying structure of the data. Specifically, a probit model may be better suited for the binary nature of the *smoker* and *exposed* variables, while for the *packs* and *hours exposed* variables, a negative binomial model may better fit the skewed nature of the data. Table 21 shows the ATEs for each of the four outcomes, for these alternative models. The results are consistent with those in the main analysis.

Table 21: Regression Model Sensitivity

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exposed** Coef./S.E.
Average Treatment Effect	0.0447	0.1055	-0.2277***	-1.3924***
	(0.0350)	(0.1208)	(0.0270)	(0.4523)
N	3299	3297	4693	4693
Sample	Ever-Smokers	Ever-Smokers	Non-Smokers	Non-Smokers
Model	Probit	Neg. Binomial	Probit	Neg. Binomial

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A further concern relates to the timing of the discontinuity. It perhaps is chance that we observe significant treatment effects at the time of the ban. To address this placebo tests, which change the timing of the discontinuity can be used to assess if significant effects are estimated when the ban did not take place. Table 22 shows results for a manipulation of the date of the ban to plus and minus one and two months. Results show that there are no significant effects at the placebo discontinuities at any of the timings, with the exception

that in the month before the ban there is a smaller drop in the probability of being exposed. These results are reassuring, but highlight there may be some anticipation of the ban in the month before it took place.

Table 22: Alternative Timing Sensitivity

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exp** Coef./S.E.
TE (-60)	0.0053	-0.0726	-0.0331	-0.0103
	(0.0392)	(0.1352)	(0.0274)	(0.5419)
N	2870	2868	4101	4101
TE (-30)	0.0067	-0.0069	-0.1071***	-0.5494
	(0.0374)	(0.1299)	(0.0273)	(0.4905)
N	3094	3092	4422	4422
TE (+-0)	0.0449	0.1076	-0.2281***	-1.3846***
	(0.0351)	(0.1220)	(0.0266)	(0.4729)
N	3299	3297	4693	4693
TE (+30)	0.0498	0.1053	-0.0327	-0.3554
	(0.0350)	(0.1200)	(0.0284)	(0.4800)
N	3117	3116	4448	4448
TE (+60)	0.0000	0.0856	-0.0017	0.2635
	(0.0370)	(0.1271)	(0.0311)	(0.4956)
N	2786	2785	3956	3956

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.8.2 Inclusion of Controls

The main model is a simple RDD model which does not include control variables. If the RDD is valid then controls should have little effect on the average treatment effects estimated. However, in the literature it is becoming more common to include controls within the model; the reason being that controls can eliminate small sample bias, improve precision and can be useful in evaluating how plausible the identification strategy is (Imbens and Lemieux, 2008). This appendix will show results from three alternative models which include controls. The first adds control variables to the main model, the second explicitly models heterogeneous treatment effects and predicts ATEs for the 'average' individual, while the third uses decomposition analysis to assess the importance of the effect of differences in covariates, rather than simply the presence of differences in covariates.

Table 23 shows results for the following model:

$$\mathbf{y} = \beta_0 + \mathbf{D}\beta_1 + \mathbf{Z}\beta_2 + (\mathbf{D}' \times \mathbf{Z})\beta_3 + \mathbf{X}_C\boldsymbol{\beta}_C + \mathbf{u} \tag{41}$$

This model is the same as the main specification (Equation 6) with a vector of controls $\mathbf{X}_C$ and corresponding coefficients $\boldsymbol{\beta}_C$ added. First, it allows for a comparison of the ATEs with those in Table 3. The ATEs and their standard errors are very similar across *packs*, *exposed* and *hours exposed*. This indicates robustness in the specification but also shows little need to add the controls in the main specification as precision is not noticeably increased.

The exception to this is *smoker*. Unlike in the main specification, where we observe no significant effect of the ban on smoking, here there is a positive effect. This implies that the null ATEs of smoker are not robust. However, further analysis shows this result itself is sensitive to the functional form of the regression discontinuity design. With each quadratic, cubic and local-linear RDD designs, when controls are included, the ATEs are not significant. Neither are there significant effects with the same specification with alternative bandwidths of 100, 50 or 30 days. This implies that the observed result is an artefact of a misspecified RDD model and random noise over the short period of time used. Indeed, when looking at Figure 2 it is clear the ban neither had an effect on smoking in the short or long run.

The additional benefit of including controls are the gains explanatory power of the model (note the R^2) and the relation of the variables to the smoking outcomes. Results show that age has little effect on smoking status, but middle aged individuals smoke more intensively and older individuals are exposed to passive smoke less. Those who have a higher SES, are female or are married are less likely to smoke, smoke less and are less likely to be exposed to second-hand smoke. City dwellers and those who have smoked in the past are more likely to be exposed and are exposed for more hours, while those who are pregnant are much less likely to smoke.

Building upon Model 41, we can identify heterogeneity not only in the levels of the smoking outcome, but in the treatment effect, by adding the interaction of the controls variables with time, the treatment dummy and the time-treatment interaction:

$$\begin{aligned}\mathbf{y} = \beta_0 + \mathbf{D}\beta_1 + \mathbf{Z}\beta_2 + (\mathbf{D}' \times \mathbf{Z})\beta_3 + \mathbf{X}_C\boldsymbol{\beta}_{C1} \\ + (\mathbf{X}_C \times \mathbf{D})\boldsymbol{\beta}_{C2} + (\mathbf{X}_C \times \mathbf{Z})\boldsymbol{\beta}_{C3} + (\mathbf{X}_C \times \mathbf{D}' \times \mathbf{Z})\boldsymbol{\beta}_{C4} + \mathbf{u}\end{aligned} \tag{42}$$

By estimating this model we estimate CATEs $\boldsymbol{\beta}_{C2}$, for each control group (in addition to the treatment effect for the reference group, β_1). This allows for the prediction of ATEs,

Table 23: RDD with Covariates

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exposed** Coef./S.E.
Treatment	0.0730**	0.1617	-0.2063***	-1.2182**
	(0.0318)	(0.1139)	(0.0256)	(0.4742)
Time	-0.0003	-0.0007	-0.0000	-0.0032
	(0.0002)	(0.0008)	(0.0002)	(0.0040)
Treatment X Time	-0.0000	-0.0005	-0.0002	0.0025
	(0.0003)	(0.0011)	(0.0003)	(0.0050)
Age	-0.0003	0.0418***	-0.0065***	-0.0799**
	(0.0026)	(0.0091)	(0.0021)	(0.0405)
Age Squared	-0.0001***	-0.0005***	0.0000	0.0003
	(0.0000)	(0.0001)	(0.0000)	(0.0004)
SES	-0.7190***	-2.7035***	-0.1870***	-5.3248***
	(0.0433)	(0.1734)	(0.0412)	(0.8709)
Sex (1 = female)	-0.0265*	-0.2425***	-0.0806***	-0.5579**
	(0.0157)	(0.0561)	(0.0130)	(0.2340)
Married Couple	-0.1106***	-0.2735***	-0.0775***	-0.3379
	(0.0173)	(0.0600)	(0.0139)	(0.2295)
Ethnicity (1 = white)	0.0047	0.4162***	0.0169	0.9995***
	(0.0364)	(0.1119)	(0.0236)	(0.3773)
Urban	0.0280	0.0571	0.0494***	0.5341**
	(0.0187)	(0.0635)	(0.0142)	(0.2529)
Pregnant	-0.2176***	-0.1534	-0.0588	-1.0966
	(0.0722)	(0.2462)	(0.0667)	(0.7064)
Ever Smoked			0.0566***	1.0001***
			(0.0128)	(0.2298)
Constant	1.0341***	1.7380***	0.8564***	8.1156***
	(0.0806)	(0.2765)	(0.0637)	(1.2810)
N	3298	3296	4691	4691
R-squared	0.1850	0.1371	0.1420	0.0381
Sample	Ever-Smokers	Ever-Smokers	Non-Smokers	Non-Smokers

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

for the 'average' individual in the sample (who has ($\mathbf{X}_C = E[\mathbf{X}_C]$) at the time of the ban. This model aims to remove selection bias from the estimates.[26] Table 24 shows the results from these predictions, where $ATE = \hat{\beta}_1 + \sum_C (E[\mathbf{X}_C]\hat{\boldsymbol{\beta}}_{C2})$. The results are consistent with those in Table 3 and 23.

An alternative approach is to use the RDD within a decomposition analysis framework. As shown in Equations (14), (15) and (16) the ATE can be decomposed into *coefficient*, *endowment* and *unexplained* effects:

[26] If, for example, there were a greater number of males just before the discontinuity then, as being female is negatively correlated with exposure, by using the standard RDD model the treatment effect would be biased downward. By estimating CATEs, then predicting the ATE for the 'average' individual then this selection bias is circumvented.

Table 24: RDD Prediction Using Conditional Average Treatment Effects

	(1) **Smoker** Coef./S.E.	(2) **Packs** Coef./S.E.	(3) **Exposed** Coef./S.E.	(4) **Hours Exposed** Coef./S.E.
Treatment	0.0701**	0.1588	-0.2020***	-1.2152**
	(0.0320)	(0.1155)	(0.0258)	(0.4859)
N	3298	3296	4691	4691

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

$$E[\mathbf{y}_1|\mathbf{Z}=\mathbf{0}] - E[\mathbf{y}_0|\mathbf{Z}=0] = \underbrace{E[\mathbf{X}_1](\boldsymbol{\beta}_1 - \boldsymbol{\beta}_0)}_{Coefficients} + \underbrace{(E[\mathbf{X}_1] - E[\mathbf{X}_0])\boldsymbol{\beta}_0}_{Endowments} + \underbrace{\beta_{10} - \beta_{00}}_{Unexplained} \tag{43}$$

In the equations that follow (unlike in Table 7 where *location exposed* variables are included) the $\mathbf{X}$ variables include only the control variables. Of interest here is the *endowment* component, $(E[\mathbf{X}_1] - E[\mathbf{X}_0])\boldsymbol{\beta}_0$, of each control variable. This shows not if there is a significant expected difference in control variables between the control and treatment groups, but if the difference is likely to contribute a significant amount to the 'total' ATE. As shown in Table 25, the *endowment* components are small and statistically insignificant for all control variables, across all four models. This provides additional assurance that it is not due to differences in covariates that we observe ATEs.

One limitation of the above approaches is that while they account for issues relating to *observable* characteristics, they do not eliminate the possibility that *unobservable* characteristics may prove problematic. However, standard RDD designs simply look to ensure there are no differences in observable control variables at the discontinuity, then assume the same to be true for unobservables. Unobservables could, therefore, be equally problematic in the the simple model. The benefit of the models proposed here is they aim to address the problems caused by observables, rather than admit defeat and not use an RDD.

Table 25: RDD Decomposed, with Covariates Alone

	(1) Smoker				(2) Packs			
	Endowments		***Coefficients***		***Endowments***		***Coefficients***	
	Coef.	S.E.	Coef.	S.E.	Coef.	S.E.	Coef.	S.E.
Time	-0.039	(0.380)	0.072	(0.283)	0.501	(1.255)	0.572	(1.008)
Controls								
Age	0.006	(0.008)	-0.357	(0.590)	0.036	(0.040)	0.062	(2.025)
Age Squared	-0.009	(0.012)	0.156	(0.307)	-0.042	(0.049)	-0.107	(1.036)
Sex (1 = female)	-0.000	(0.001)	0.012	(0.034)	0.002	(0.003)	-0.014	(0.122)
Married Couple	-0.002	(0.003)	-0.010	(0.041)	-0.004	(0.005)	-0.221	(0.149)
Ethnicity (1 = white)	-0.001	(0.001)	-0.045	(0.138)	0.002	(0.003)	0.242	(0.410)
Urban	-0.002	(0.002)	-0.062	(0.064)	-0.004	(0.007)	0.034	(0.223)
Pregnant	0.000	(0.001)	-0.003	(0.004)	0.001	(0.003)	-0.001	(0.014)
Ever Smoked	0.000	(.)	0.000	(.)	0.000	(.)	0.000	(.)
Unexplained			0.385	(0.320)			0.183	(1.094)
Component Effects	-0.009	(0.006)	-0.308	(0.319)	-0.010	(0.015)	-0.003	(1.108)
Total Effect	0.068*	(0.034)			0.169	(0.120)		
Time Controls	YES				YES			
N	3298				3296			

	(1) Exposed				(2) Hours Exposed			
	Endowments		***Coefficients***		***Endowments***		***Coefficients***	
	Coef.	S.E.	Coef.	S.E.	Coef.	S.E.	Coef.	S.E.
Time	0.367	(0.290)	-0.286	(0.211)	-3.897	(6.039)	3.492	(5.145)
Controls								
Age	-0.001	(0.003)	-0.392	(0.446)	0.012	(0.037)	-9.815	(8.520)
Age Squared	-0.001	(0.002)	0.358	(0.229)	-0.014	(0.047)	5.446	(4.328)
Sex (1 = female)	-0.000	(0.001)	0.030	(0.029)	-0.001	(0.006)	-0.120	(0.561)
Married Couple	0.001	(0.001)	0.020	(0.035)	0.003	(0.008)	0.638	(0.513)
Ethnicity (1 = white)	0.002	(0.002)	-0.192**	(0.088)	0.015	(0.022)	-0.125	(1.178)
Urban	0.000	(0.000)	0.067	(0.045)	0.006	(0.011)	0.544	(0.885)
Pregnant	-0.000	(0.001)	0.001	(0.003)	-0.004	(0.008)	0.028	(0.029)
Ever Smoked	-0.000	(0.000)	-0.008	(0.026)	0.000	(0.002)	0.315	(0.498)
Unexplained			-0.089	(0.228)			1.921	(4.316)
Component Effects	-0.000	(0.006)	-0.115	(0.227)	0.017	(0.033)	-3.089	(4.316)
Total Effect	-0.204***	(0.026)			-1.151*	(0.484)		
Time Controls	YES				YES			
N	4691				4691			

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.9 Conditional Average Treatment Effects: Sensitivity

Within the main analysis, we estimate Conditional Average Treatment Effects (CATEs). These effects are estimated conditional on age-specific socioeconomic status (an index we derive), using our approach, which uses the Beta distribution to derive weights within a linear regression discontinuity framework. To identify how sensitive the estimates for CATEs are, this appendix will show results for different conditional variables, estimation methods and regression discontinuity designs.

A.9.1 Socioeconomic Status Sensitivity

To identify if effects are sensitive to the choice of conditional variable five alternative variables are used within separate estimations. The five alternatives are socioeconomic status (SES), log equivalised income (income), the Index of Multiple Deprivation (IMD), the National Statistics Socio-economic Classification (NS-SEC) and the highest level of qualification. SES: Alternative is similar to our SES variable; it is estimated using principal component analysis using the same variables (shown in Table 1), however it is estimated for the whole population, not separately for each age group. The log equivalised income is derived in the survey, IMD is discretised into quantiles (by the survey), while the NS-SEC and highest qualification variables are discrete categories as shown in Table 1. To allow comparability between discrete and continuous variables, for the discrete variables the graphs shown use quantiles within that discrete category, resulting in the horizontal size of each category corresponding to the proportion of the sample within that category. The quantiles are ordered to have the 'lowest' category at the bottom; so NS-SEC categories are ordered from *Never Employed* to *Higher Manager/Professional*, while highest qualification begins at *None* and ends with *Degree*.

For both *active* smoking variables, smoker and packs, the non-significance of the SES: Age results hold across the five alternative measures; so for conciseness these results are not shown. Results for *passive* smoking variables, exposed and hours exposed, are shown in Figure 10 and 11, respectively. In Figure 10 the magnitude, significance and nature of heterogeneity appear very similar across all alternative variables. Effects for those at the lowest quantiles are either small or insignificant, while the effects become stronger as the quantile increases. The exception to this is that those with the highest qualifications how smaller effects than those with A Levels. The effects also appear slightly higher for those at the higher quantiles of income, compare to those higher quantiles of SES. In Figure 10 we observe similar effects for SES, SES: Alternative, income and IMD. However, the effects are

less similar for NS-SEC and highest qualification. In general, we observe that the precision of the estimates is lower for the discrete categories, in comparison to the continuous variables.

Figure 10: Conditional Average Treatment Effects: Exposed

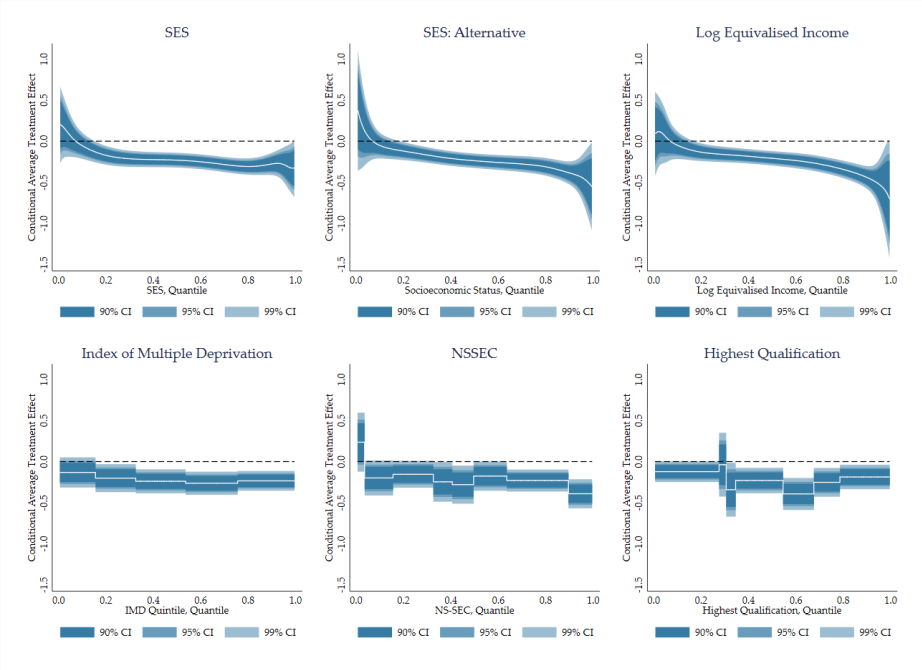

Figure 11: Conditional Average Treatment Effects: Hours Exposed

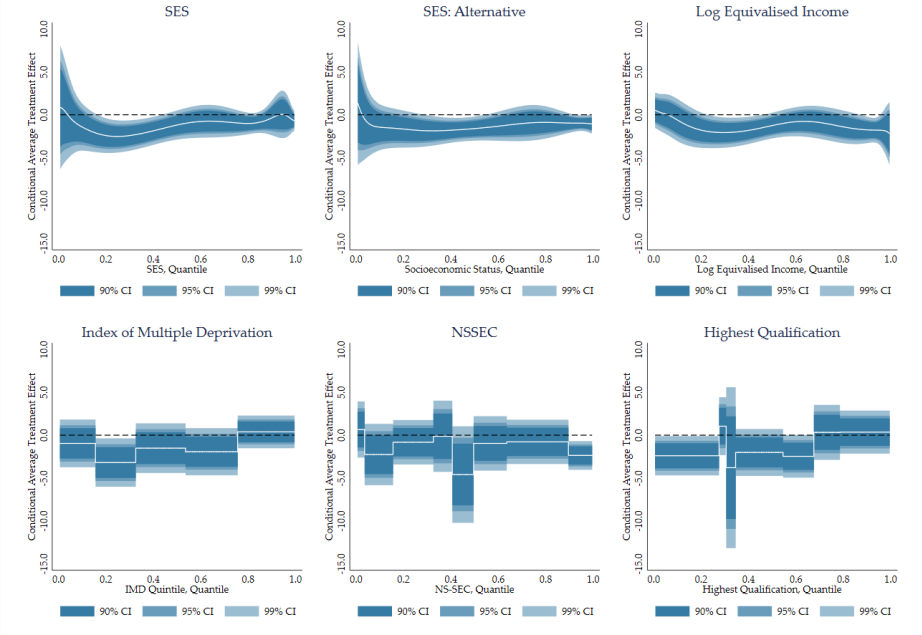

A.9.2 Method Sensitivity

The main analysis estimates CATEs using our method which uses the Beta distribution to derive weights for each individual. However, alternative methods can be used to estimate these effects. We will focus on the *hours exposed* outcome to illustrate these different methods. First, two entirely separate methods: dummy interaction models, with varying number of interactions, and parametric interaction models, with varying degrees of freedom. Second, within the non-parametric framework estimations using different weighting functions (Beta: Constant, Uniform, Triangular, Normal and Epanechnikov) and different degrees of precision, s, within the Beta weighting function.

First, Figure 15 shows the comparisons of two, five, ten, twenty and fifty dummy interactions with our method. This method allows conditional effects to be non-parametric, as they are separately estimated for groups within particular quantiles. The trade-off with the dummy interaction model is that for low numbers of interactions the underlying heterogeneity may be missed, while for high number of interactions precision of the estimates is much reduced. Results do, however, seem consistent with our method, however, our method is able to capture heterogeneity, with relatively high precision, and with a smooth curve.

Figure 12: Conditional Average Treatment Effects: Dummy Interaction

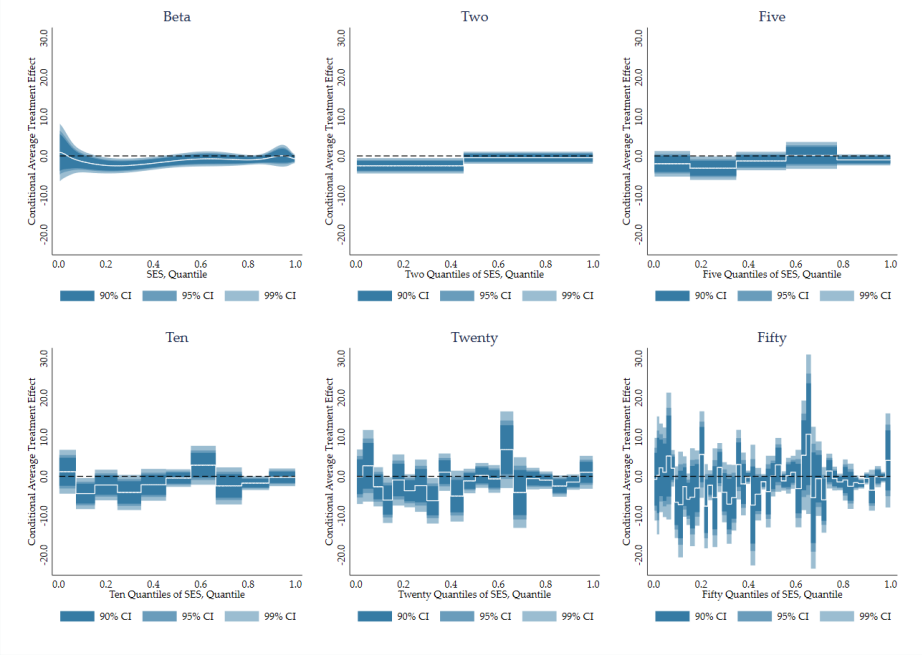

Figure 13 shows an alternative method, interacting standard RDD model with socioeconomic status. The base case, no interaction (the same as the average treatment effect) is compared to linear, quadratic and cubic interaction terms. Each additional polynomial degree adds flexibility to the estimation. The cubic model closely resembles the non-parametric estimations of our model, while each lesser degree is removes the flexibility to estimate as accurately. We do, however, observe that precision of the estimates decreases as the higher degree terms are added.

Figure 13: Conditional Average Treatment Effects: Interaction

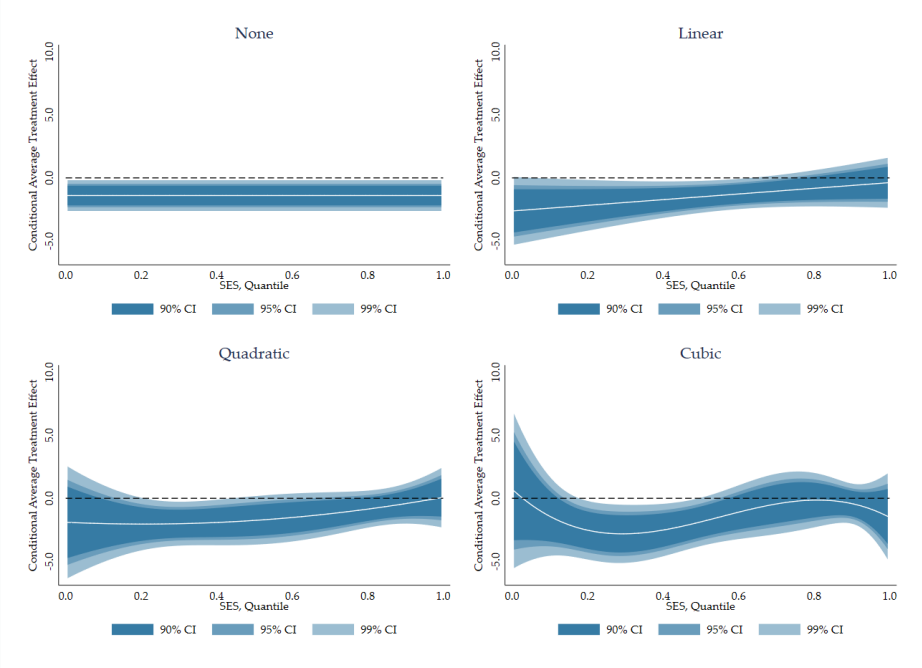

When we estimate CATEs a particular weighting function, for each local regression, has to be assumed. We derive equations for and use the Beta distribution as the weighting function within our main model. To test if the estimated effects are sensitive to this choice of weighting function we estimate results for the Beta: Constant, Uniform, Triangular, Normal and Epanechnikov. The latter four are common weighting functions, used within the kernel estimation and local-linear regression literature, while the first uses an alternative assumption relating to variance (see Appendix A.3). Results in Figure 14 show that the weighting function used has little impact of the shape of the CATE curve. It does, however highlight several differences. The uniform, and to some extent triangular, weighting functions give less smooth estimates. Other weighting functions also struggle to estimate effects close to the bounds. Here, we plot the CATEs against the expected weighted rank $(\sum_i^N (w_{ji} r_i) / \sum_i^N (w_i))$

of each quantile, which is the empirical quantile of the conditional variable. The Beta function is most capable of approaching the bounds (with the expected weighted rank). However, as observed it is still prone to loss of precision close to the bounds.

Figure 14: Conditional Average Treatment Effects: Different Weights

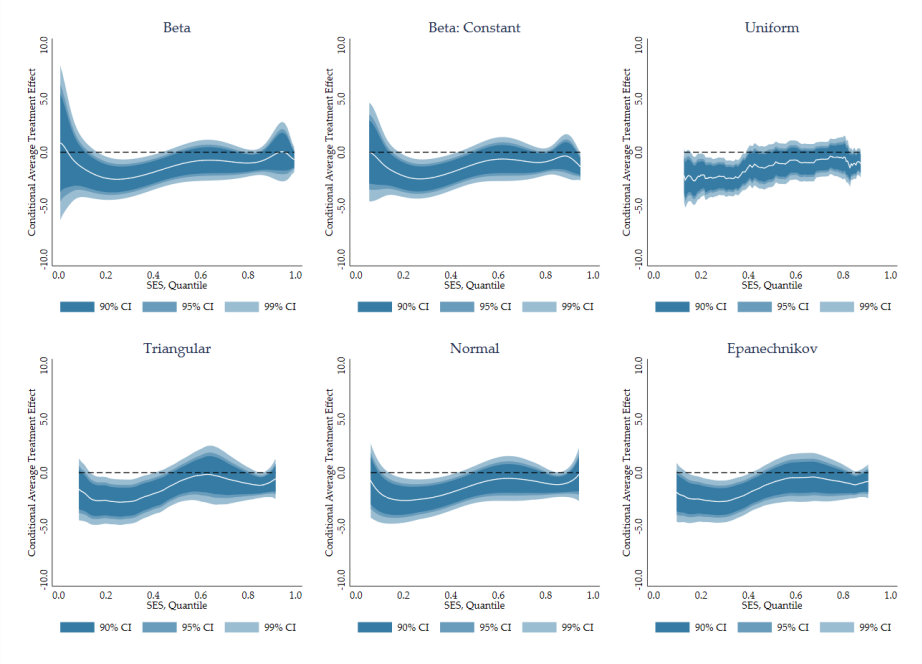

Given that the Beta distribution is assumed as the weighting function the precision parameter, s, still has to be chosen. This choice of parameter is common within the non-parametric literature, within the above other mentioned weighting functions the 'bandwidth' similarly has to be chosen. Figure 15 illustrates the consequences of the choice. A trade-off between smooth curves with tighter confidence intervals, for lower precision parameters, and estimates with less bias, for higher precision parameters, emerges. Higher precision parameters give much lower weight to those far away from the quantile of interest, meaning less data is effectively used in each estimation, while lower precision parameters give more uniform weights, hence providing more power. Regardless of the precision parameter we do, however, observe a similar trend in the CATEs.

While alternative methods, weighting functions and precision parameters can be used to estimate CATEs, the results estimated here seem relatively robust to these alternatives.

Figure 15: Conditional Average Treatment Effects: Precision

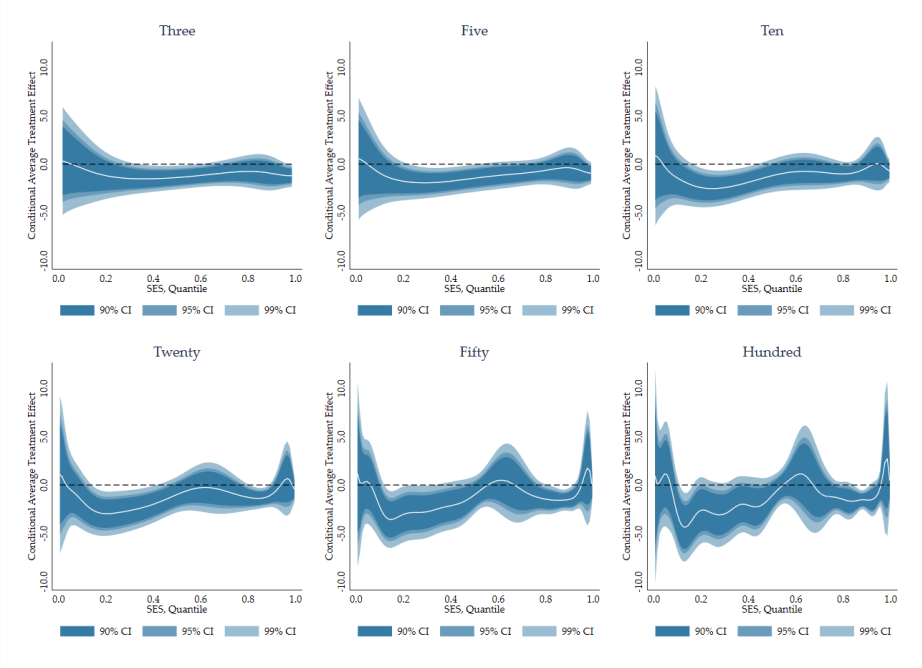

A.9.3 RDD Sensitivity

Within the main estimation the regression discontinuity design assumes linearity across the forcing variable, time. To test how sensitive the estimates CATEs are to this assumption Figure 16 show results where no time trend is assumed and Figure 17 where a quadratic term is used.

In general, the results are similar. There is no significant effect for smoker or packs across the majority of socioeconomic status (SES) and significant reductions for both exposed and hours exposed. For exposed, the size of the effect increases with SES, while for hours exposed we observe larger effects are for those who are poor (but not extremely poor) compared to those with the highest SES. What is clear, however, is the results for the poorest are sensitive to the functional form assumed. The reason for this is a greater heterogeneity amongst the very poor, meaning that average levels of the outcome, in each time period, are noisier. The estimate of the time trend is, therefore, sensitive to the functional form assumed. More observations, more often, would however address this issue.

Figure 16: Conditional Average Treatment Effects: No Time

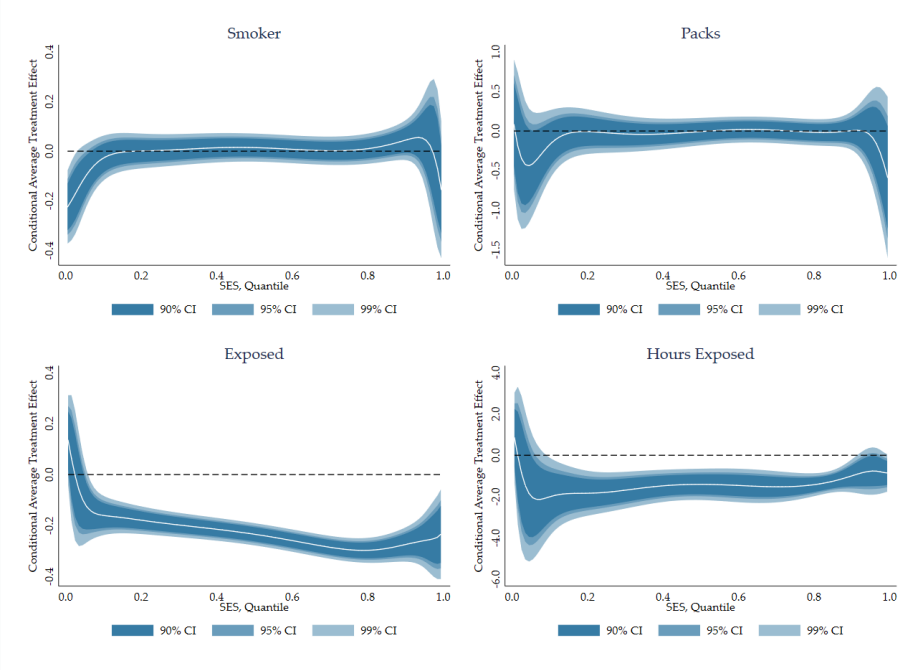

Figure 17: Conditional Average Treatment Effects: Quadratic Time