European University Institute

DEPARTMENT OF ECONOMICS

Essays in Applied Economics

Alastair Ball

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Florence, 03 December 2015

European University Institute
Department of Economics

Essays in Applied Economics

Alastair Ball

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Examining Board

Prof. Jérôme Adda, Bocconi University and EUI, Supervisor
Prof. Andrea Ichino, EUI
Prof. Christian Dustmann, University College London
Prof. Petter Lundborg, Lund University

Abstract

This thesis considers two critical periods in life that can have long-term effects on health and prosperity. The first paper provides new evidence on the consequences of foetal exposure to high levels of pollution for the risk of stillbirth, and for the long-term health and labour market outcomes of those that survive. Variation in in utero exposure comes from a persistent weather system that affected London for five days in December 1952, preventing the dispersion of atmospheric pollution. This increased levels of total suspended particulate matter by around 300%. Unaffected counties in England are used in a difference-in-differences design to identify the short and long-term effects.

Historical registrar data for the nine months following the smog show a 2% increase in reported stillbirths in London relative to national trends. As foetal deaths often go unreported, the exercise is then repeated for registered births. The data show around 400 fewer live births than expected in London, or a reduction of 3% against national trends. Survivors are then identified by district and quarter of birth, and their health and labour market outcomes observed at fifty and sixty years old. Differences-in-differences estimates show that survivors are in general less healthy, less likely to have a formal qualification, and less likely to be employed than those unaffected by the smog.

The second chapter considers the decision over which skills to acquire at university - taken at seventeen, this decision has significant impacts on both unemployment on graduation, and long-term incomes over the life cycle. Under the hypothesis that more expensive tuition might lead students to acquire skills in high demand, this paper examines the effects of the 2006 increase in fees from 1,200 to 3,000 a year on the probability that a given student would study a stem subject. A propensity matching methodology is used to control for sample selection caused by reduced university participation following the increase in fees. Results indicate that the fee increase caused a five percent reduction in the probability that a given student would study a stem subject. Course level data from the Higher Education Statistics Agency suggests that that subjects most affected were nursing, pharmacology, and medical technology.

Acknowledgements

I am grateful to a great number of people for support and advice during the PhD - but first, I'd like to thank my family for all their unwavering support and patience for the last few years and in particular my parents, who raised me with the helpful belief that all avenues were open to me, and the knowledge that they would support me whatever came. The EUI has been a really enriching place to spend half a decade, and I'd like to thank my old undergraduate supervisor, David Cohham, for good advice when I was at a fork in the road. Goodness knows how life would have gone in other programmes, but I can't imagine having had such a good and interesting time.

I was very lucky in my academic advisors, and I'd like to express my gratitude to Jerome Adda and Andrea Ichino for enduring my wide range of research plans with patience and humour, and for giving great advice when it counted. I learned so much through all the different generations of the 'Adda group' about new techniques, strategies for research, and generally how to talk about and engage in economics. In addition to Jerome, I'd like to thank all the other members of the group, who all taught me so much. I gained a great deal from all the taught material we encountered, and I'd like to particularly say thanks for everyone involved in our Causality course, which was probably about the most interesting course I've encountered; and in Micro III, - my own personal Kobayashi Maru. I'd also like to thank Jessica, Julia, Marcia and Lucia for all their help throughout, and of course Thomas for his tireless help in overcoming obstacles erected by the ONS, Eurostat and others. Lastly, I'd like to thank Juan Dolado for invaluable help in my final year of the PhD.

Of course, ninety percent of my time was spent with my friends and colleagues at the EUI, in my cohort and others. I really couldn't imagine a better bunch of people to spend time with. A special thanks to Claudia and to flatmates Abian, Andrew, Gama, and Robert for making my time in Florence such a fun and interesting experience.

Contents

Chapter 1

Air pollution, foetal mortality, and long-term health: Evidence from the Great London Smog

I would like to thank Jerome Adda, Andrea Ichino, David Levine, Juan Dolado, Brandon Restrepo, Claudia Cerrone, Nicola Rogers, Rachel Stuchbury, Chris Marshall and the participants in the EUI job market seminar for helpful comments during the development of this paper. The permission of the Office for National Statistics to use the Longitudinal Study is gratefully acknowledged, as is the help provided by staff of the Centre for Longitudinal Study Information & User Support (CeLSIUS). CeLSIUS is supported by the ESRC Census of Population Programme (Award Ref: ES/K000365/1). The authors alone are responsible for the interpretation of the data. This work contains statistical data from ONS which is Crown Copyright. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

1.1 Foetal exposure to air pollution

The effects of exposure *in utero* are a largely hidden cost of pollution, but the impact on health and well-being can be profound. The goal of this paper is to study the short and long-term effects of foetal exposure to a strong pollution episode - the London Smog of 1952. This persistent smog was caused by a high-pressure weather system that hung over the city for five days, preventing the normal dispersion of pollution. During these days, visibility dropped to just a few metres, and ambient levels of pollution increased by up to 300%. There was no panic at the time as Londoners were accustomed to winter fogs, but the health effects were later found to be severe, with around 12,000 excess deaths eventually attributed to the smog. The contemporaneous impact of the smog on adult health can be seen in Figure 1 below, showing quarterly deaths in London as a proportion of those registered in England and Wales.

There are a number of reasons why the London smog is a good environment in which to study the effects of foetal exposure to pollution. First, the unusually sharp exposure makes it possible to observe health effects by trimester, and to differentiate cleanly between foetal and neonatal effects. This is unusual as most plausibly exogenous variation in pollution, such as variation from recessions

Figure 1.1: Ratio of deaths during the smog: London, and England & Wales

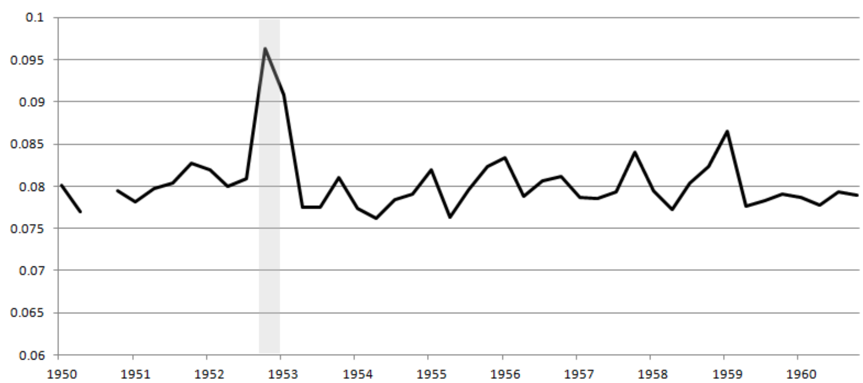

Notes: Shaded area represents the fourth quarter of 1952 - the period of the London Smog. Data on mortality from the English Registrar General. Q3 1950 is missing from the dataset.

or policy changes, take one or two years to take full effect[1]. Second, as the smog occurred in 1952, it is possible to observe both the contemporaneous effect on foetal mortality, and the health and labour market outcomes of survivors over the subsequent fifty or sixty years. Lastly, levels of pollution were particularly high during the period. This is important because the majority of research on the health effects of pollution comes from high-income countries, where information on pollution and health are of good quality, but where pollution levels are generally low. While pollution levels[2] in developed countries are typically lower than $30\mu g/m^3$, those found in most low- and middle-income countries are generally much higher. In 2012, over a hundred cities had average levels of over $100\mu g/m^3$. Average levels in Delhi were $286\mu g/m^3$, while those in Peshwar, home to over three million people, were over $500\mu g/m^3$. As the shape of the dose-response of health to foetal exposure to pollution is still a subject of ongoing research, the London smog presents a unique opportunity to learn about the long-term health effects of the severe pollution levels currently found in much of the world[3].

Foetal deaths often go unreported in official statistics, and can even go unnoticed by the mother in some circumstances. As a result, the existing evidence on prenatal effects often exploit indirect measures of foetal survival. Jayachandran (2009) uses variation in atmospheric pollution from a large wildfire that swept through Indonesia in 1997. The wildfires caused severe atmospheric pollution, with levels of PM10 in some areas exceeding 1000 $\mu g/m^3$ for several days during the fires. The identification strategy exploits geographic variation in pollution and the sharp timing of the wildfire to test for reductions in the number of live births caused by prenatal mortality. Results show that the pollution led to 15,600 missing children, an effect driven predominantly by *in utero* exposure. Sanders and Stoecker (2011) take another indirect approach, exploiting the Trivers-Willard hypothesis that male foetuses are more vulnerable to external shocks than females. Using plausibly exogenous variation from the 1970 Clean Air Act and data from a 50% sample of U.S. birth certificates, the authors find

[1] Two examples of research based on 'sharp' events are Currie and Walker (2011), who observe the effects of reduced traffic congestion near highways caused by the introduction of an automated tolling system; and Jayachandran (2009), who studies the health effects of wildfires that swept through Indonesia in 1997. Both papers are discussed in more detail below.

[2] All levels of pollution refer to annual average levels of particulate matter less than ten microns across(PM10). Source for data: the World Health Organization, available at *www.who.int/topics/airpollution/en*

[3] For an overview of the dose-response of health to atmospheric pollution, see Zivin & Neidell (2013). For a discussion of cross-country differences between dose-responses, see Arceo-Gomez et al (2014).

that a one standard deviation drop in annual average particulates reduces the percentage of male live births by 3.1%. Results from both papers suggest that foetal survival is strongly linked to the education and income levels of the mother.

A possible issue in observing the health effects of foetal exposure to pollution among those surviving till birth is sample selection. If stillbirth caused by an *in utero* shock is more likely for less healthy foetuses, those observed after birth would be selected from the strongest in the cohort[4]. Bozzoli, Deaton and Qintana-Domeque (2009) encounter this effect when studying the cross-country link between child mortality and adult height. The authors find the expected relationship among most countries, but find that child mortality is associated with an *increase* in the height of the surviving population in the poorest countries, where mortality was particularly high. There is, nonetheless, strong evidence that foetal exposure to pollution affects infant health and mortality[5]. Chay and Greenstone (2003) use variation in pollution caused by the 1981-82 U.S. recession to study the effects of particulate pollution on infant mortality. They find that a 1% decrease in pollution in a county results in a 0.35% reduction in the infant mortality rate. The strongest effects were found for infants less than one month old, suggesting that foetal exposure was an important factor. Currie and Walker (2011) study the effects of air pollution caused by traffic using variation from the introduction of the EZ-Pass scheme in New Jersey and Pennsylvania. This scheme allowed drivers to pass through toll gates without stopping, and resulted in a sharp reduction in carbon monoxide pollution in residential areas close to the tolls. The paper uses a differences-in-differences strategy, comparing the change in infant health of those *in utero* close to highway tolls to those *in utero* close to other parts of the highway system. Their results show that the introduction of the EZ-Pass scheme resulted in around an 11% reduction in prematurity, and a 12% reduction in birth weight - a common proxy of infant health.

There is much less evidence from low and middle-income countries, due mostly to the difficulty in obtaining information on health and pollution. Arceo-Gomez, Hanna & Oliva (2014) gather ten years of weekly data on health for forty eight municipalities in Mexico City, where data for pollution is also available. The authors adopt an IV strategy using temperature inversions - which prevent the dispersion of atmospheric pollution - as an instrument for exposure. The IV estimates show that a $1000\mu g/m^3$ increase in particulates results in 0.24 infant deaths per 100,000 births - a health effect similar to those found in the literature on the United States. Greenstone and Hanna (2014) construct a database of infant health and pollution levels in India in order to study the effectiveness of environmental regulations. The authors also test the effects of the most successful of the reforms, which promoted the use of catalytic converters, on infant mortality. Their results were suggestive of a decline in infant mortality, but were not statistically significant.

Very few papers study the long-term effects of foetal exposure to pollution, mainly because of the difficulty in obtaining information on place of birth for individuals observed as adults. Sanders (2012) overcomes this issue with administrative education data from Texas that contains information on both students' test scores in high school and their counties of birth. Exposure to pollution is calculated using county-level data on total suspended particulate matter, and is instrumented using county-level changes in relative manufacturing employment. The author finds that a standard deviation decrease in particulates is associated with a 2% increase in grades using OLS, and in a 6% increase using IV. Isen, Rossin-Slater and Walker (2014) use linked administrative data from the U.S. census to investigate the effects of foetal exposure to particulates on incomes later in life. To

[4]See Almond and Currie (2011) for a fuller discussion of 'culling' in papers studying the effects of foetal shocks.

[5]See Zivin & Neidell (2013) for a survey of the health effects of pollution, including the effects of foetal exposure on infant mortality. Currie and Vogl (2013) provide an overview of the long-term effects of early shocks in developing countries, including those from atmospheric pollution.

identify the effect, the authors exploit a sharp drop in atmospheric pollution that followed the implementation of the 1970 Clean Air Act. Their results indicate that a 10 unit decrease in particulates resulted in a 1% increase in earnings for individuals aged 29-30, mostly driven by a drop in labour force participation.

This paper takes a differences-in-differences approach to measure the short and long-term health effects of foetal exposure to pollution, comparing outcomes in London over time to those in unaffected counties of England and Wales. In order to study the effects of foetal exposure to pollution on prenatal mortality, quarterly historical registrar data on stillbirths is gathered for the period from 1948 to 1964. Differences-in-differences estimates on stillbirths in the period from one to nine months after the smog show a 2% increase against national trends. As foetal deaths are often unreported, the exercise was repeated for numbers of live births. The data show around 400 fewer live births than expected, or a reduction of 3% against national trends. Overall, the results indicate that *in utero* exposure to air pollution had a strong effect on foetal mortality, much of which went unreported in official statistics.

Survivors are then identified in the ONS Longitudinal Study - a 1% sample of the population of England and Wales - using information on district and quarter of birth. Two related designs are used to observe the long-term effects of *in utero* exposure to pollution. The first is a differences-in-differences design comparing outcomes for those conceived in 1952 (and affected by the smog during the pregnancy) to those conceived in 1953, using unaffected counties of England and Wales to control for year-level effects. The second design splits London into 'low pollution' districts that experienced weaker (but still very high) pollution during the period, and 'high pollution' districts that experienced especially severe pollution. The estimates from this design compare outcomes for children conceived in 1952 in the high or low polluted areas with outcomes for children conceived in 1953 in the same areas, again using unaffected counties of England and Wales as a control.

As the short-term effects indicate that neonatal mortality was affected by the smog, the first outcome observed in the ONS-LS was the gender of survivors when they were first observed in the 1971 census. In contrast to other papers in the literature, there is little evidence that female foetuses were more likely to survive. Simple time series estimates of the proportion of males in the affected cohort in London are, in general, not significantly different from zero. There is weak evidence for the opposite effect - that males were more likely to survive than girls. Estimates on the proportions of males among those affected in the first and second trimesters, when children are especially vulnerable, are positive. An estimate of the effect for the areas of London most seriously affected are also positive, with a 95% confidence interval of [-0.01, 0.10]. These results are not supportive of the Trivers-Willard hypothesis that female foetuses are more likely to survive an adverse shock, but do not contradict previous studies finding this effect. The pollution levels observed during the smog are far higher than those in the United States, where most work on the Trivers-Willard hypothesis has been conducted. It is interesting to note that survivors affected as infants are more likely to be female, suggesting that there may be differences in the effects on survival of shocks experienced *in utero* and as a newborn.

The analysis of long-term effects focuses on individuals observed in 2001 and 2011 (at around fifty and sixty years old) in the ONS-LS. However, linked information on deaths makes it possible to observe mortality in youth and middle age. Results show that those affected by the smog are, on average, two percent less likely to die before sixty than their peers. The effect is statistically significant and appears strongest for those aged over forty five. Two hypotheses might explain this counter-intuitive result. The first is that 'what doesn't kill you makes you stronger', with a maternal or a foetal response to the shock protecting the child in later life. The second is that, as with Bozzoli et al (2009), those surviving the early health shock were drawn from the strongest in the cohort.

Results from other outcomes are strongly supportive of the latter hypothesis: those observed in later life were 2% more likely to report themselves in poor health, 3% less likely to have an A-level qualification, and 1% less likely to be employed at fifty then their peers. The employment effects are driven almost entirely by males, who were 4% less likely to be employed at fifty than their peers.

Differences in gender balance among those first observed in the ONS-LS in 1971 suggest that the effects on foetal survival were much less severe in the 'low pollution' districts. As a result, estimates from these districts are likely to incorporate a smaller bias towards health caused by strong-survivor effects. Individuals born in these districts were 6% more likely to report themselves in poor heath, 10% less likely to have an A-Level qualification, and 4% less likely than their peers to be employed at fifty. In the 'high pollution' districts, where the gender balance among survivors shifted 4% towards males, the estimated differences between survivors and their peers were weaker, less likely to be statistically significant, and even positive in some circumstances. Individuals from these districts were 2% less likely to have an A-Level qualification, but were 3% *less* likely than the unaffected population to report themselves in poor health, a pattern observed in both male and female survivors. If the individual-level health effects in the 'high pollution' districts were as strong, or stronger than those in the 'low pollution' districts (where the impact was still over $800\mu g/m^3$) then foetal mortality must have had a profound impact on the characteristics of the surviving population. It should be noted that the estimated effects on the health and labour market outcomes of those in the 'low pollution' districts, although strong, are also likely to be lower-bounds of the true health effects for individuals.

This paper contributes to the literature in a number of ways. Most evidence on the effects of foetal exposure to pollution is based on relatively low levels of pollution - this paper provides new evidence on the effects of foetal exposure to severe pollution of a kind closer to those currently experienced in middle and low income countries. In the two areas of London studied, particulate pollution increased by an average of $800\mu g/m^3$ and $1800\mu g/m^3$ during the five days of the smog. Data from daily readings in Delhi[6], where the average level of pollution is currently around $280\mu g/m^3$, show how variable pollution levels can be. Between 2004 and 2010, there were fifty six occasions in which the Town Hall pollution meter recorded particulate levels over $800\mu g/m^3$, and nine occasions when the levels were over $1000\mu g/m^3$. The Town Hall meter provides observations for only around one day in ten during the six years observed, but it is clear that the high annual levels of pollution recorded by the World Health Organisation are likely to hide a large number of severe pollution shocks of the kind observed in London 1952. Second, the unusually sharp variation in pollution caused by the smog makes it possible to separate effects by trimester and to differentiate cleanly between effects from prenatal and neonatal exposure. Third, the high quality data collected at the time makes it possible to gather evidence on foetal mortality, which is often under-reported, or impossible to observe at the levels of pollution studied. Fourth, the long period since the smog makes it possible to observe the long term effects of foetal exposure for up to sixty years, for a variety of outcomes.

Taken as a whole, the results suggest that the London smog had a significant impact on foetal mortality, and limited the health, investment in education, and employment prospects of those that survived. Restricting attention to just mortality, the World Health Organisation estimates that pollution caused 3.7 million premature deaths worldwide in 2012, 88% of which occurred in low and middle income countries. The results of this study suggest that this figure is an underestimate, missing both the deaths, and the long-term scarring of those not yet born.

[6] Available from the Indian Central Pollution Control Board.

1.2 The London smog of 1952

On the fifth of December of 1952, winds dropped, and a high pressure weather system settled over London. Atmospheric pollution from traffic and the burning of coal that was normally dispersed by convection became trapped by the resulting temperature inversion, and a thick ground-level smog formed over the city. The weather conditions persisted for five more days, during which time visibility dropped to metres, and pollution levels increased threefold. London was accustomed to 'pea-soupers' - thick winter fogs - and there was little panic. The people of London stayed at home, which was the official advice at the time, and normal life continued on the eleventh. When official figures on deaths and hospitalisations arrived a week later, it became clear that something quite serious had happened. A report by the Ministry of Health (1954) attributed over four thousand deaths[7] to the smog, leading Parliament to pass the Clean Air Act of 1956, drafted with the goal of preventing any further smogs in London.

Figure 1.2: Pollution in December 1952 - By district

Notes: Points marked on the map indicate the locations of the twelve PM10 monitoring stations active during the smog. Light shaded areas indicate an average increase of around $800\mu g/m^3$, dark shaded areas indicate an average increase of $1800\mu g/m^3$. Pollution data from the Fuel Research Board and Wilson (1954), mapping data from the Ordinance Survey.

There is good information available on pollution levels during the smog. Figures 1.3, 1.4 and 1.5 show daily measurements taken during the first half of December for London, Great Britain, and for other large towns. The December smog appears not to have affected rest of Britain: there is a small increase in pollution levels in the other big towns on the seventh and eighth of December, but

[7] Later studies have revised this number up to 12,000 Bell & Davis (2001)

nothing close to the scale of the London smog. The smog affected rich and poor areas alike - with Kensington, Chelsea and South London among the worst hit areas. Figure 1.2 shows the effects of the smog by London borough. The Ministry of Health (1954) report divided London into two levels of impact using a combination of measurements of sulphur dioxide and particulates. The darkest areas in the map are the most seriously affected, experiencing an increase of around $1800\mu g/m^3$. The lighter gray areas indicate areas experiencing lower - but still very high - increases of around $800\mu g/m^3$. Average levels of pollution during the period from 1949/50 to 1954/5[8] can be seen in figure 1.6. For the period for which there is data available, there is an upward trend in average levels of Sulphur Dioxide, with a dip in the year prior to the smog.

Figure 1.6: Annual levels of sulphur dioxide, parts per million

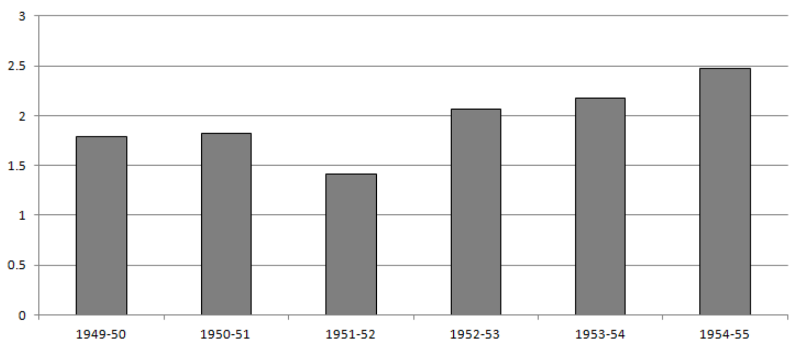

Notes: Reported years start in October (at the start of the winter peak.)
Pollution data from the Fuel Research Board.

There is less information available[9] for England & Wales, but figures 1.17 and 1.18 in the appendix show information on emissions of smoke and sulphur dioxide and - as cold weather prevents the dispersion of atmospheric pollution - on monthly minimum temperatures. Figure 1.17 shows an increase in the emission of sulphur dioxide, similar to the trend for London shown in figure 1.6. Figure 1.18 shows monthly minimum temperatures for London and England for the period. Winter temperatures in London during 1952 were low relative to long-term trends, but almost identical to those in the five year period from 1951 to 1955. With the exception of the smog in December, average levels of pollution in 1952 appear broadly consistent with the general trends in London and the rest of the UK.

The four thousand excess deaths recorded in the next three months were initially attributed to influenza, but there was no evidence of influenza in the lungs of the diseased[10], and the Chief Medical Officer concluded that there was no major outbreak of influenza in 1952[11]. The majority of those that died during the smog were over forty five years old. Figure 1.7 shows the total number of deaths recorded during the weeks following the smog, broken down by age. Figure 1.1, shown earlier in the paper, shows the ratio of deaths in London to those in England and Wales. The impact of the smog is clearly visible, and there are no comparable incidents in the ten year period that the data covers.

[8]Each yearly observation begins in October: the Fuel Research Board reported averages this way to avoid splitting the winter peak into two observations.

[9]Pollution data from before 1960 is not stored centrally: the information presented comes from the records of the Fuel Research Board that were stored in the National Archives when the Board was disbanded. Only a given percentage of records were kept, making the construction of long series challenging. The gap in the quality of the information available for December 1952 and other times is due to the Fuel Research Board's particular interest in this episode.

[10]Ministry of Health (1954)

[11]Bell and Davis (2001)

Figure 1.3: Pollution in December 1952 - London

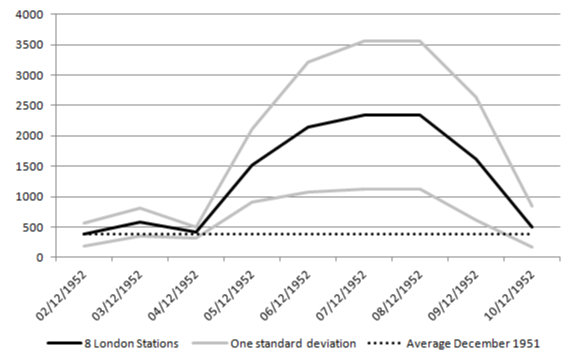

Notes: Pollution data is for PM10 particulates and is measured in microgrammes per cubic metre. Solid black line based on an average from 8 monitoring stations in London. Gray line shows the standard deviation of this average. Dotted line shows the December average from 1951. Data from the Fuel Research Board

Figure 1.4: Pollution in December 1952 - Great Britain (excluding London)

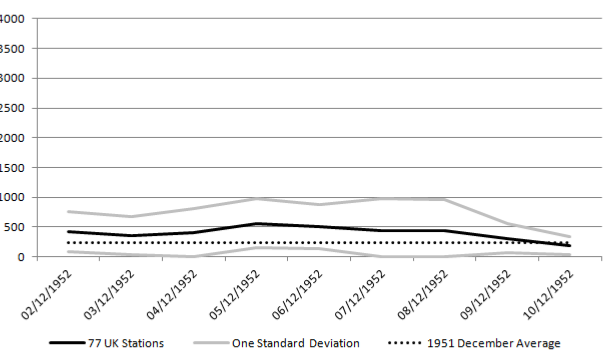

Notes: Pollution data is for PM10 particulates and is measured in microgrammes per cubic metre. Solid black line based on an average from 77 monitoring stations in Great Britain. The gray line shows the standard deviation of this average. Dotted line shows the December average from 1951.Data from the Fuel Research Board

Figure 1.5: Pollution in December 1952 - Other big towns (excluding London)

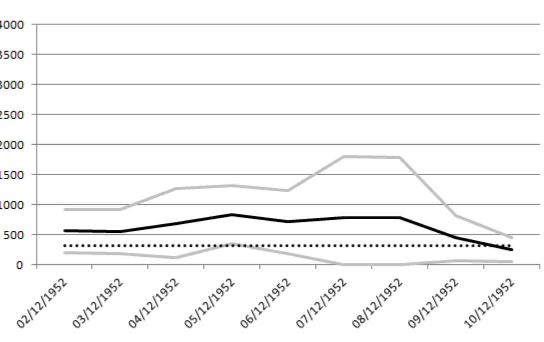

Notes: Pollution data is for PM10 particulates and is measured in microgrammes per cubic metre. Solid black line based on an average from 16 monitoring stations other large towns in the UK. These are: Bradford, Bristol, Cardiff, Glasgow, Leeds, Leicester, Liverpool, Manchester, Newcastle, Nottingham, and Sheffield. The gray line shows the standard deviation of this average. Dotted line shows the December average from 1951.Data from the Fuel Research Board

Figure 1.7: Deaths in London during the smog of 1952, by age

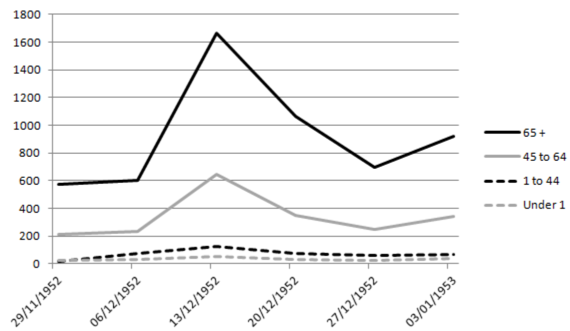

Notes: Deaths reported at the end of each week to the London administrative county, recorded in Ministry of Health (1954)

1.3 The effect of the smog on foetal mortality

This section studies the effects of the London smog of 1952 on foetal mortality. Foetal mortality often goes unreported, and can even go unnoticed when it occurs early in the pregnancy. For this reason, an analysis of how stillbirths were affected by the smog will be supplemented with an analysis of the effects on live births. As discussed in Jayachandran (2009), evidence on 'missing children' can be a good proxy for foetal mortality, especially when official statistics are likely to provide an incomplete picture. Data on stillbirths and births comes from the Quarterly Report of the Registrar General, which form the basis for official statistics in the UK on births and deaths.

Stillbirths in London and in England & Wales can be seen in figure 1.8 below. The series for London is relatively volatile, but appears to share both quarterly fluctuations and the effects of nation-wide shocks with the English series. Figure 1.9 plots the ratio of these series. The first three quarters of 1953 are highlighted - this is the period in which stillbirths among children affected by the London smog would be observed. Although there appears to be an increase in relative reports of stillbirths, it is not clear whether this is a true effect or natural variation. In order to determine whether the effect is statistically significant, the series for London and England & Wales are normalised by population[12], and the effect of 1953 tested using the differences-in-differences model below.

$$\Delta S_t = \alpha + \beta 1953_t + \gamma \phi(t) + \boldsymbol{\lambda} quarter_t + \epsilon_t \tag{1.1}$$

The left hand side is the difference in stillbirths in London and England & Wales per one thousand people, ϕ is a polynomial function of time used to capture any secular trends, *quarter* is a vector of dummies to control for seasonal effects, and 1953_t is a dummy taking the value of one for 1953, when stillbirths caused by the smog would be observed. Estimates for β, which should capture any effect on stillbirths, are shown in the table below, for polynomial trends of different orders. Estimated coefficients for all specifications are positive, with a central estimate of 0.006. This figure is equivalent to an extra 20.2 stillbirths in 1953, or an increase of just over 2%.

The effect on live births

A complementary approach for testing for foetal mortality is to observe births over time, and to check for 'missing' children. The strength of this approach is that, as stillbirths are often unreported, it can

[12] A more common normalisation for stillbirth data is against total births - in this case this would be inappropriate as a foetal insult would be expected to affect both series.

Figure 1.8: Stillbirths in London (left axis), England & Wales (right axis)

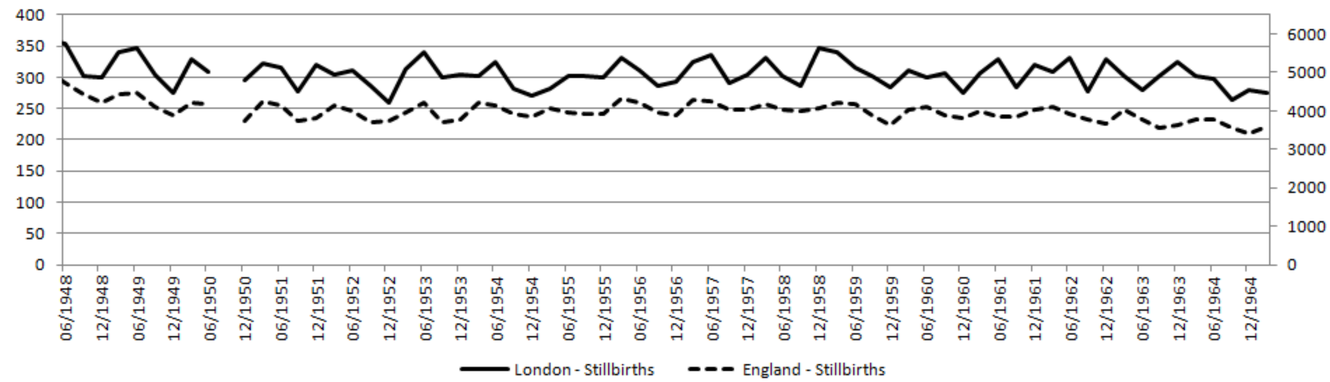

Notes: Data from the Registrar General

Figure 1.9: Ratio of stillbirths in London, and England & Wales

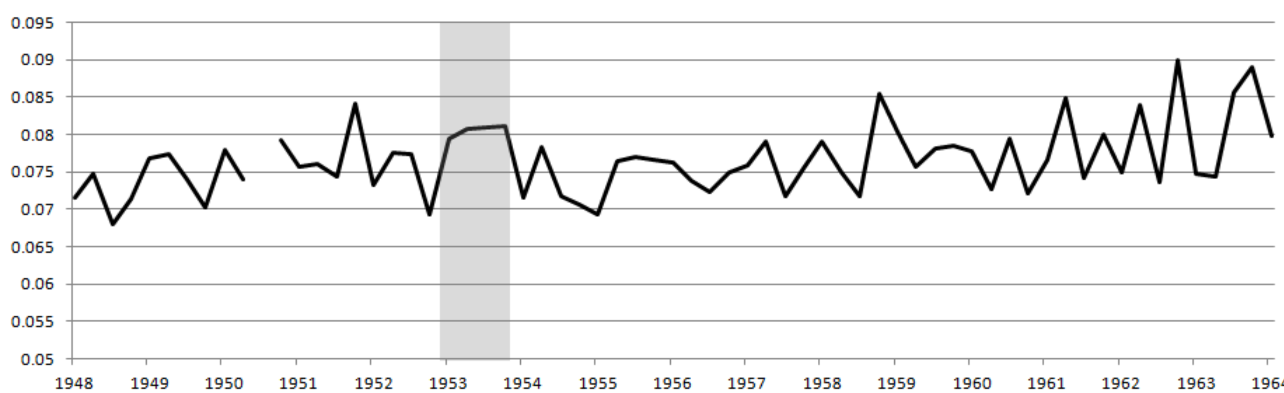

Notes: Data from the Registrar General. Shaded area highlights where increased stillbirths due to London smog would be observed.

$\hat{\beta}$	Standard Error	Polynomial Trend	Adjusted R^2
0.006*	0.002	2nd Order	0.68
0.006*	0.003	3rd Order	0.68
0.006*	0.003	4th Order	0.70
0.006*	0.003	5th Order	0.69
0.006*	0.003	6th Order	0.69

β measures the difference in stillbirths per 1000 people in London and in England & Wales. Estimated coefficient is equivalent to an increase of just over 2%. Stars indicate significance at the 5%, 1% and 0.1% levels

give a more accurate picture than official figures. Figure 1.10 shows births over time in London and in England & Wales. As with stillbirths, the two series share seasonal and broad long-term trends. Figure 1.11 shows the ratio of the two series. The plot is less volatile than that for stillbirths, and appears to show two well defined events. The first is the 1952 London smog - there appears to be a clear drop in the quarters in which foetuses affected by the smog could have been born. The second is the London smog of 1948, which was less severe than that of 1952, but still resulted in high levels of pollution. Those affected in the first and second trimester, when foetuses are most vulnerable, would have been born in the third and second quarters of 1949 respectively, where there appears to be a sharp drop in births in London. The secular trend in the series is more complex than that for stillbirths, with a increase, a decrease, and a levelling out of the relative numbers of births in London compared to England & Wales.

The statistical significance of this drop in births was tested in the same manner as with stillbirths. The results can be seen in the table below. The model with the fifth order polynomial is the most parsimonious model that fits the data well. As seen in figure 1.11, the secular trend has highly non-linear shape that lower order polynomials have trouble matching. Adding a sixth term does not improve the overall fit of the model or change the estimates substantively. The figure of -0.128 is equivalent to 404 fewer births than expected, or a drop of 3%. This estimate is higher than for stillbirths, implying that much of the prenatal mortality caused by the smog went unrecorded.

1.4 Long-term effects for survivors

The previous section showed that *in utero* exposure to severe pollution can have a significant impact on prenatal mortality: this section studies the impact on those that survive. Information on long-term outcomes comes from the U.K. Office of National Statistics Longitudinal Study. This is a survey based on a 1% sample of the decennial census for England and Wales, and has been linked to data on major events such as births, migrations, and deaths[13]. The survey holds information on around a million individuals, with around 423,000 linked to both the 2001 and 2011 censuses. Membership of

[13]The ONS-LS has also been linked to widow or widower status, and cancer registrations.

Figure 1.10: Births in London (left axis), England & Wales (right axis)

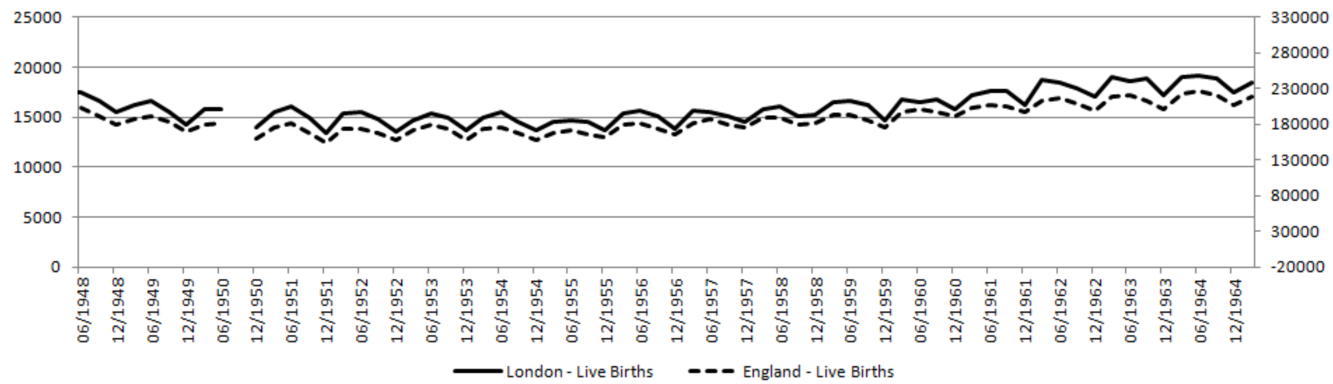

Notes: Data from the Registrar General

Figure 1.11: Ratio of births in London, and England & Wales

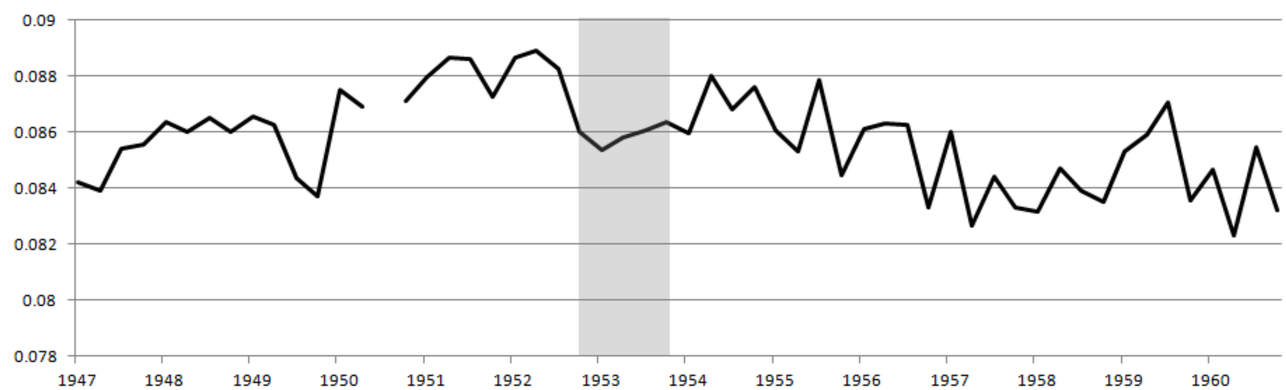

Notes: Data from the Registrar General

$\hat{\beta}$	Standard Error	Polynomial Trend	Adjusted R^2
-0.061	0.054	2nd Order	0.91
-0.078	0.055	3rd Order	0.92
-0.069	0.048	4th Order	0.93
-0.124^{*}	0.043	5th Order	0.95
-0.128^{*}	0.043	6th Order	0.95

β measures the difference in births per 1000 people in London and in England & Wales. Estimated coefficient is equivalent to an increase of just over 3%. Stars indicate significance at the 5%, 1% and 0.1% levels

the survey is determined by being born on one of four dates in a year.

Individuals affected by the smog *in utero* are identified by their quarter and district of birth. The basic identification strategy is to compare the outcomes of the London '1952 Cohort' that were conceived in 1952 and exposed to the smog in December with outcomes for the '1953 Cohort' conceived in the subsequent year. The main comparison is illustrated in the figure below. An advantage of this design is that, aside from personal exposure to the smog, the two cohorts are exposed to most aggregate level shocks at very similar ages[14]. Looking forward, the two cohorts enter the labour market at almost the same time. Looking back, any non-health related effects from the smog, such as issues created by deaths in the family would be expected to affect the two cohorts in a similar way. Notice that individuals conceived in the first quarter are excluded from both cohorts. This is because it is not possible to determine whether these individuals were exposed to the smog shortly before, or shortly after birth.

As was seen in figure 1.6, the average pollution levels during pregnancy are very similar for the 1952 and 1953 cohorts. As was seen in figure 1.18 in the appendix, weather in the two years was also similar. However, in order to control for any potential year-level effects, a differences-in-differences strategy is employed, using 1952 and 1953 cohorts from other counties of England & Wales, which were unaffected by the smog. The design does not study those *in utero* during 1951 (the year before the smog) because these children would have been exposed to the smog as one year olds, an age at which the children are still very vulnerable. Later years were also rejected - from the data available average levels of pollution appear appreciably higher[15].

The main equation to be estimated for an individual i born at time t in county c is:

$$Y_{ict} = \alpha + \boldsymbol{\beta}\mathbf{County}_c + \delta D_t^{1952} + \gamma D_c^L D_t^{1952} + \epsilon_{ict} \tag{1.2}$$

Where Y_{ict} is the outcome variable, $\mathbf{County}_c$ is a vector of county-level dummies, D^{1952} is a dummy for those conceived in 1952, and D^L is a dummy taking the value one for those born in London. The

[14] Overall employment levels when the two cohorts entered the labour market can be seen in figure 1.19 in the appendix. The (affected) 1952 cohort entered the labour market between 1969 and 1971, when unemployment was 2.8% on average. The (unaffected) 1953 cohort entered between 1970 and 1972, when unemployment was 3.1% on average.

[15] An earlier version of this analysis did include this extra year, comparing outcomes for the 1952 cohort to those for both the 1953 and 1954 cohorts. Results are the same.

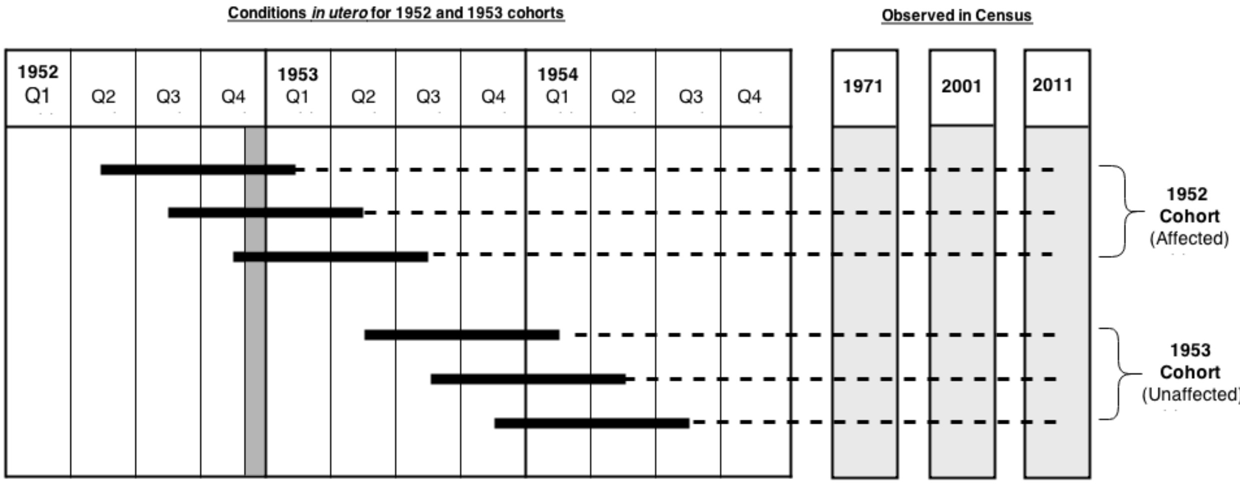

Notes: Thick lines represent time spent *in utero*. Dotted line marks progression through the years after birth. The shaded area denotes the London smog of December 1952. Members of both cohorts are observed in the ONS-LS in 1971, 1981, 1991, 2001, and 2011. The analysis on the gender of survivors focuses on information from the 1971 census (when individuals were around 19.) The analysis on health and other outcomes focuses on information from individuals observed in 2001 and 2011, when they were around fifty and sixty years old.

effect of the smog on outcome Y_{ict} should be summarised in γ. A second specification incorporating information on the severity of exposure to pollution is also estimated. Boroughs of London are divided into 'high' and 'low' levels of exposure using the Ministry of Health (1954) schema, illustrated in figure 1.2. In this design, the basic comparison is of outcomes for those *in utero* in high or low polluted areas with outcomes for those *in utero* in 1953 in the same area. Unaffected counties outside of London are again used to control for year level effects.

$$Y_{ict} = \alpha + \boldsymbol{\beta}\mathbf{Area}_c + \delta D_t^{1952} + \gamma_L D_c^{Low} D_t^{1952} + \gamma_L D_c^{High} D_t^{1952} + \epsilon_{ict} \quad (1.3)$$

This specification is identical to equation 2, except that the **County** vector of dummies has been replaced with **Area**, as London's county has now been split into two parts. The dummy D^{Low} takes a value of one if individual i was *in utero* in a part of London less affected by the smog. The dummy D^{High} works in the same way for more severely affected parts of London. The parameters of interest are γ_L and γ_H, giving the effects of the smog on outcome Y in the areas experiencing low and high exposure during the smog.

Table 1.1 reports summary statistics for the four groups studied in the design. The first column shows information on individuals *in utero* in London during the smog. Compared to individuals in the subsequent year, they are slightly more likely to be male, are less likely to have either A-level or degree qualifications, are more likely to be in poor health, and are less likely to be employed. Differences are small for most measures. The third and fourth column show figures for the 1952 and 1953 cohorts born in unaffected parts of the country. In limiting illness and employment, they show the same small increase, meaning that England-wide trends might be driving some of the effects seen in London.

Table 1.1: Summary statistics for the 1952 and 1953 cohorts

	London		England & Wales	
	1952 Cohort	1953 Cohort	1952 Cohort	1953 Cohort
Sample size	850	840	4040	4060
Male	0.50	0.49	0.50	0.51
A-levels	0.28	0.32	0.26	0.26
Degree	0.20	0.24	0.20	0.20
Poor health	0.07	0.06	0.09	0.09
Limiting illness	0.12	0.11	0.16	0.15
Employed	0.85	0.86	0.82	0.83

Notes: Percentages for poor health, limiting illness and employment recorded at age fifty. Sample sizes are rounded to the nearest ten. Source: ONS Longitudinal Survey.

Gender

As there is evidence that exposure to the smog resulted in foetal mortality, a natural first step is to gather evidence on who survived. The approach follows that of Sanders and Stoecker (2011), who use shifts in the gender of survivors as evidence of culling, motivated by the Trivers-Willard hypothesis that female foetuses may be more likely to survive adverse shocks. In this case, shifts in the gender composition will be taken as evidence of when exposure is most harmful, and how the effect on foetal mortality differed by intensity of exposure.

Figure 1.12: Percentage of males in ONS-LS sample

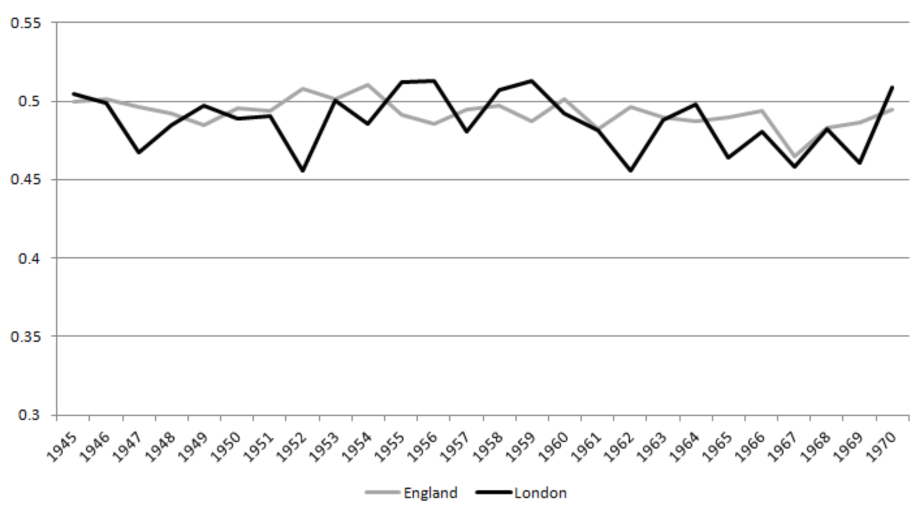

Notes: Source: Office of National Statistics Longitudinal Study

The analysis focuses on individuals aged around 19, when they were first observed in a representative sample of the 1971 census for England and Wales. Figure 1.11 shows the percentage of males in the ONS-LS sample, by year of birth. Both the series for London and England & Wales are relatively volatile, but there is a clear drop in the proportion of males for those born in 1952[16]. These individuals would have been under eleven months old when affected by the smog. There may also be signs of variation during those that experienced the 1962 and 1948 smogs as infants. There is no evidence, however, of variation among those that would have been *in utero* during the smog, and born in 1953.

In order to study the different effects of neonatal and prenatal exposure to the smog, the following equation was estimated separately for London and the rest of England & Wales using the full series of birth dates.

$$Prob(Male)_{it} = \alpha + \beta year_t + \boldsymbol{\delta}\mathbf{quarter}_t + \gamma\mathbf{impact}_t + \epsilon_{it} \tag{1.4}$$

Where *year* is a simple linear trend, **quarter** is a vector of dummies to control for season, and **impact** is a vector of dummies for people born in the quarters before and after the smog. Results are summarised in the table below. For London, only one coefficient is significant at the 5% level - the analysis is based on some small sample sizes, particularly at the trimester level in London, and all other coefficients are imprecisely estimated. However, the central estimates for those affected during the smog are not supportive of the Trivers-Willard hypothesis that females are more likely to survive a foetal shock. For foetuses affected in the first and second trimester, when they are expected to be most vulnerable, the estimates point in the opposite direction.

Table 1.2: Proportion of males, by age at exposure to smog

Quarter of Birth	Age affected	London	s.e.	England & Wales	s.e.
1952 Q1	11 months	-0.03	0.033	-0.01	0.014
1952 Q2	8 months	-0.09^{*}	0.034	-0.01	0.014
1952 Q3	5 months	0.02	0.031	0.02	0.014
1952 Q4	**2 months**	**-0.02**	**0.033**	**0.02**	**0.014**
1953 Q1	3rd Trimester	-0.03	0.033	0.01	0.014
1953 Q2	2nd Trimester	0.04	0.030	0.01	0.014
1953 Q3	1st Trimester	0.03	0.030	0.00	0.013

Notes: Sample size for London is 29830. The sample size for England & Wales is 153880. Age during the smog is correct to within one month.
Source: Office of National Statistics Longitudinal Study

In order to observe prenatal effects by intensity of exposure to pollution, equation (4) is simplified by replacing the vector **impact** with a simple dummy for the 1952 cohort that was exposed to the smog *in utero*. Estimated coefficients for England & Wales, London, and the districts of London affected by 'high' and 'low' levels of pollution during the smog can be seen in the table below. In London and the districts that were less affected by the smog, there is no evidence of any effect. In areas of London that were most severely affected, the coefficient suggests a 4% upwards shift in the proportion of males. The estimate is not significant at standard levels, but the 95% confidence interval is [-0.01, 0.10]. Overall, there is no strong evidence of gender-biased survival for those affected *in utero*. In the context of this severe foetal shock, there is no support for the Trivers-Willard hypothesis. Most estimates, though imprecisely estimated, suggest that it was male, rather than female, foetuses that were more likely to survive the smog. Although not the focus of this paper, it is interesting to note that females affected by the smog after birth did appear more likely to survive, suggesting that the effects of pre- and neonatal exposures to atmospheric pollution may be quite different.

[16]Observing sample sizes by gender shows the same story, and suggests that this effect is driven entirely by a drop in males.

Table 1.3: Proportion of males by strength of exposure

England & Wales	London Low exposure	London All districts	London High exposure
0.00	0.00	0.01	0.04†
(0.007)	(0.020)	0.016	0.026

Notes: † indicates significance at the 10% level. 95% confidence interval is [-0.01, 0.10]. Sample size for London is 29830. The sample size for England & Wales is 153880. For high and low districts, 10770 and 13460. Age during the smog is correct to within one month. Source: Office of National Statistics Longitudinal Study

Mortality

There is little evidence to suggest what the effect of foetal exposure to pollution on the long-term mortality of survivors might be. On one hand, there is evidence that those surviving are less likely to perform well in school[17] and less likely to be in employment[18], suggesting a negative effect on health. On the other hand, evidence in the literature and from the earlier analysis show that foetal exposure has a significant effect on survival till birth. If survivors are drawn from the strongest in the cohort, they may be more likely to survive than cohorts in which less healthy individuals are still observed in the data. Figure 1.13 shows the cumulative hazard of death for people born in London, comparing outcomes for the affected 1952 cohort, and the unaffected 1953 cohort. The cumulative hazard of death for the unaffected cohort appears to accelerate at the age of around forty five, relative to that for the affected cohort. There may also be signs that the unaffected were more likely to die in their early twenties, but the effect is smaller.

Figure 1.13: Cumulative hazard of mortality - London

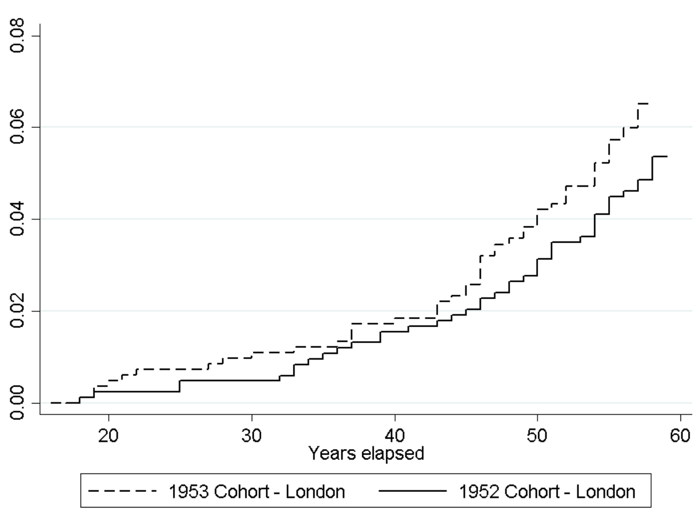

Notes: Nelson-Aalen cumulative hazard function. Source: Office of National Statistics Longitudinal Study

[17] Sanders (2012)
[18] Isen, Rossin-Slater and Walker (2014)

Differences in differences estimates are presented in table 1.4. An important factor in interpreting these estimates is the assumption of common trends: fitting a linear trend shows that older individuals born in London and the rest of England & Wales are more likely to die before 2011. The difference in the trends for London and England & Wales is 0.0005 which, although significantly different from zero at the 5% level, is too small to seriously affect the estimates. Two measures of mortality are reported. The first is the probability of dying before 2011, when the two cohorts were around sixty years old. Estimates come from a simple linear probability model (OLS), and results can be interpreted as the percentage change in the probability of dying before 2011 caused by *in utero* exposure to the smog. The result for individuals born in all of London confirm that individuals affected were 2% less likely to die than their peers. The results for those born in areas of London affected by high and low levels of pollution from the smog are 2% and 3% respectively. All estimates are significant at the 0.1% level. The second measure is from a Cox proportional hazards model. This divides the hazard of dying in a particular year into a baseline hazard, determined by time alone, and a component that is affected by a given set of covariates. 'Ties' in the data are handed using the Breslow method. Estimates show a positive effect on the baseline hazard when they are greater than one, and a negative effect when they are less than one. The results from the two methods are essentially identical.

Education

Educational attainment can capture the effects of weaker health problems that might not appear in hospital records and mortality statistics. There is already evidence from Sanders (2012), who studies pupils in Texas, that foetal exposure to pollution can affect educational attainment later in life. The pollution levels studied in this paper are much stronger than those in Texas; it is possible that foetal mortality could bias estimates towards health (and better educational attainment) by the removal of weaker individuals. Nonetheless, results show that those affected by smog *in utero* are less likely than their peers to hold formal qualifications.

Figure 1.14 shows educational attainment by birth year for those born in London and the rest of England & Wales. Until 1954, the probability of holding a qualification in both England & Wales and London increase by around a percentage point every five years, and then level off. Two asymmetric shocks can be seen for London. The first is a dip in 1949 - this is the cohort that would have been *in utero* during the winter smog of 1948. This event was less severe than that in 1952, but was still a very serious pollution episode. The second is the Great London Smog of 1952. There is a noticeable dip for two cohorts - those born in 1952, who were affected by the smog as infants, and those born in 1953, who were *in utero* during the smog.

Table 1.4: Pollution and Mortality: all London, and two levels within London

	Died		Hazard of Death	
	London	Levels	London	Levels
1952 cohort	0.00	0.00	1.14	1.13
	(0.005)	(0.005)	(0.100)	(0.097)
London	0.00	-	1.00	-
	(0.003)		(0.059)	
Lond. High	-	0.00	-	1.00
		(0.003)		(0.061)
Lond. Low	-	0.00	-	1.06
		(0.003)		(0.061)
Smog impact	-0.02^{***}	-	0.72^{***}	-
	(0.005)		(0.003)	
High impact	-	-0.02^{***}	-	0.74^{***}
		(0.003)		(0.063)
Low impact	-	-0.03^{***}	-	0.63^{***}
		(0.003)		(0.063)
N	9720	9720	9720	9720
R^2	0.00	0.00	-	-

Notes: 1952 cohort refers to those *in utero* during the smog. Estimated smog impact from the interaction of treatment area and membership of 1952 cohort. Standard errors clustered at county level: stars indicate significance at 5 and 1, and 0.1%. Results from OLS linear probability model indicate percentage increase in probability of dying in sample. Estimates from the Cox proportional hazard model indicate a positive effect on the hazard of dying in a given year when above one, and a negative effect when below.

Figure 1.14: Percentage with formal qualifications, by year of birth

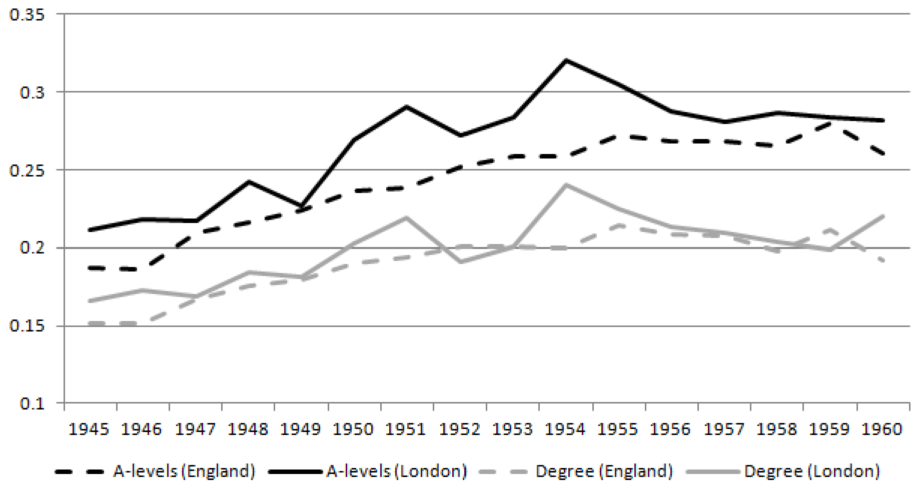

Notes: In the London series, those born in 1949 *in utero* during smog of 1948. Those born in 1952 affected as infants by the 1952 smog. Those born in 1953 affected *in utero* by the smog. Source: Office of National Statistics Longitudinal Study

Although the overall England & Wales and London trends are similar until 1954, the assumption of common trends is stronger in this case than with the other outcomes studied. A test of these trends over the sample supports this - with educational attainment in London growing around 0.2 of a percentage point faster in London than in England & Wales. The estimated effects of the smog are between fifteen and forty times larger, however, and are unlikely to be driven by this difference.

Differences-in-differences results can be seen in table 1.5 below. Two outcomes are studied. The first is the probability of holding an A-level qualification. This is a secondary level qualification taken at the age of seventeen or eighteen. The second is the probability of holding a degree. Both are binary variables, and are estimated using a linear probability model (OLS). Estimated effects can be interpreted as the change in the probability of holding a qualification with a change in the independent variable. Results from the main specification show that survivors of *in utero* exposure to the smog are 3% less likely to hold an A-level and 5% less likely to hold a degree than their peers. Separating London according to the severity of pollution exposure during the smog, survivors from the the 'low' treated area are far less likely to hold qualifications than those from the 'high' treated area. One explanation for this effect is the bias towards health (and educational attainment) caused by weaker individuals dying before being observed. Recall that those in the 'high' pollution area also seemed to have the largest in gender, implying that the effects on foetal mortality were strongest here.

Health

Individuals in the sample were asked to rate their health at the age of fifty (into three categories) and sixty (into five categories). Although self-reported health is a widely used measure in the health economics literature, it may not capture all health effects. Deaton (2008) discusses three key issues. The first is that people might not fully perceive the impacts of a health shock. Someone with small respiratory problems may not fully contemplate the career as a professional footballer they might have had in full health. The second is that people grow accustomed to their ailments, and no longer

Table 1.5: Educational Attainment: all London, and two levels within London

	A-levels		Degree	
	London	Levels	All London	Levels
1952 cohort	0.00	0.00	0.00	0.00
	(0.013)	(0.013)	(0.012)	(0.011)
London	0.06^{***}	-	0.04^{***}	-
	(0.011)		(0.009)	
Lond. High	-	0.05^{***}	-	0.02^{*}
		(0.011)		(0.009)
Lond. Low	-	0.07^{***}	-	0.06^{***}
		(0.010)		(0.009)
Smog impact	-0.03^{*}	-	-0.05^{***}	-
	(0.013)		(0.012)	
High impact	-	$-0.02^{\dagger}$	-	-0.01
		(0.013)		(0.012)
Low impact	-	-0.09^{***}	-	-0.10^{***}
		(0.013)		(0.012)
N	6830	6830	6830	6830
R^2	0.00	0.00	0.00	0.00

Notes: 1952 cohort refers to those *in utero* during the smog. Estimated smog impact from the interaction of treatment area and membership of 1952 cohort. Standard errors clustered at county level: stars indicate significance at 5 and 1, and 0.1%. Results from OLS linear probability model indicate percentage increase in probability of holding the relevant qualification. † Significant at the 10% level - 95% confidence interval: [-0.05, 0.00] consider them to be day-to-day problems. The third is that there are cross-country differences in how this kind of question is answered due to both cultural differences, and differences in the average health of comparison groups. In the context of this study, the first two problems might result in an underestimate of any health effects. The issue of cross-country comparisons is unlikely to be important in this study as all people sampled are from the same, relatively homogeneous country. A second measure of heath will also be briefly discussed: the response to the question ‘Do you have a condition that limits daily activities?’. Although this question is relatively concrete, it would also result in an underestimate of true health effects if unhealthy individuals had less strenuous ‘daily activities.’

Figure 1.15 shows responses to questions about health in 2001 and 2011, by year of birth. As might be expected, those that are younger when asked the question are in better health. In general[19] the London and England & Wales series follow each other closely. In both England & Wales and London there is a slight dip, against trend, in the early fifties. Those affected as infants would have been born in 1952, while those affected *in utero* would have been born in 1953. There are no obvious effects of the smog in either series.

Figure 1.15: Poor health in 2001 (right axis) and 2011 (left axis), by year of birth

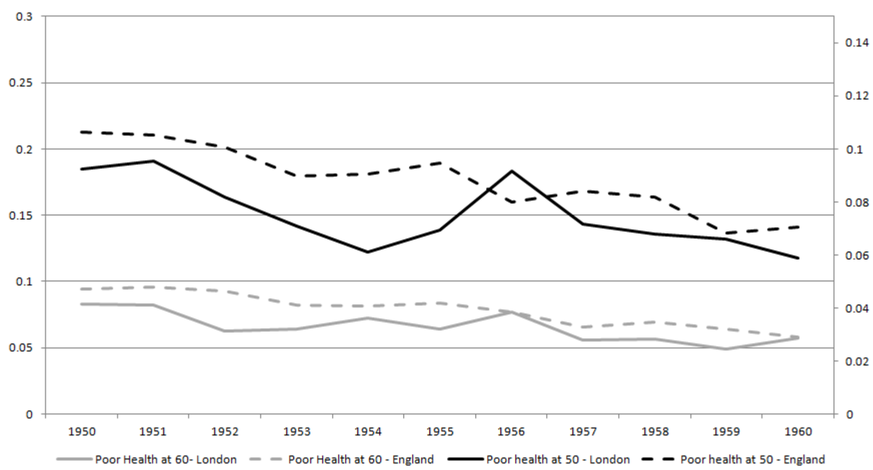

Notes: ‘Poor health‘ at fifty defined by lowest response to a three category question on current health. ‘Poor health‘ at sixty defined by lowest two responses to a five category question on current health. Source: Office of National Statistics Longitudinal Study

The difference in trends between England & Wales and London is statistically significant but small - at around 0.3 of a percent[20] at fifty and 0.2 of a percent for poor health at sixty. In principle, this difference (which would make those affected by the smog appear less healthy) could drive results. In practice, this is unlikely as results from the differences-in-differences model are precisely estimated, and an order of magnitude larger.

Table 1.6 shows differences-in-differences estimates from the linear probability model (OLS). The health measure used is a binary variable taking the value one if a person reports themself as being

[19]The exception is the health in 2001 for people born in 1956. There does not appear to have been a large event in London that might explain this data point. (The Asian flu epidemic was a year later and would also have affected England & Wales.)

[20]This is somewhat driven by the outlier in 1956 - without it, the difference is under 0.2.

in poor health[21]. The categorical health variable answered in 2001, when the sample of interest are around fifty, has three categories. The question answered at sixty has five categories, with the 'poor health' dummy taking a value of one for the two bottom categories. Consequently, the direction, but not the magnitude of estimated effects are comparable between the 'aged fifty' and 'aged sixty' results.

Results for all districts of London show that those surviving the *in utero* exposure to pollution are 2% more likely to report themselves as being in poor health. By sixty, there is no health difference between those affected and their unaffected peers. These estimates are for the same individuals observed ten years apart: as those affected *in utero* are unlikely to have improved in health between fifty and sixty, this result can most easily be explained by the unaffected group 'catching up' in terms of poor heath by age sixty.

Dividing London into areas affected by high and low levels of pollution during the smog reveals significant heterogeneity in estimates. Those in the area affected by 'low' pollution show more serious health effects than the London average, at both fifty and sixty. Those that were in the area affected by 'high' pollution - where foetal mortality appears to have been most significant - are healthier than their unaffected peers. The effects are large and statistically significant, with individuals 3% less likely to be in poor health at fifty, and 4% less likely at sixty.

Effects on limiting illness

The exercise was repeated for a second measure of health - a binary measure for people declaring a health problem that limits their daily activities. Results are complementary to those for the measure of poor health discussed above. Those affected by the smog *in utero* were 1% more likely to have a limiting illness at both fifty and sixty, but the coefficient was not precisely estimated. Dividing London into areas affected by high and low levels of pollution reveals the same effect as with poor health: those observed at fifty in the low area were 3% more likely to report a limiting illness, while those observed in the high area were 3% less likely to report a limiting illness. Both estimates are significant at the 5% level.

Employment

Foetal exposure to pollution could affect employment directly, through its effects on health, or indirectly through its effects on educational attainment. As with education, changes in employment can be a good measure of the kind of health effects that would not appear in statistics on hospitalisations or mortality. The studied cohorts entered the labour market under very similar conditions. Figure 1.19 in the appendix shows that average levels of unemployment on entry were 2.8% while those for the 1953 cohort were 3.1%, and so any employment effects are unlikely to be driven by issues of timing. This section studies employment at the age of fifty[22]. Figure 1.16 shows the proportion of people employed by year of birth. As with the other series, London and England & Wales have similar trends, with the exception of those born in London in 1956. During the period from 1950 to 1954, the trends in the two series are essentially parallel, but a relative dip can be seen for those born in 1952 (and exposed to the smog as infants.) There is no clear effect for those born in 1953, who would have been affected *in utero*.

[21] Results from the full categorical health variable are identical in terms of overall direction of effects and significance, and can be found in the appendix. Estimated coefficients from the ordered logit are in terms of percentage changes to log odds, which do not have a particularly intuitive interpretation.

[22] Employment at sixty is not as good a measure of health effects: early retirement could either signal a successful career or poor health, and the distinction is not clear in the data.

Table 1.6: Poor health: all London, and two levels within London

	Aged 50		Aged 60	
	London	Levels	All London	Levels
1952 cohort	0.00	0.00	0.00	0.00
	(0.007)	(0.007)	(0.006)	(0.008)
London	-0.03^{***}	-	-0.01	-
	(0.006)		(0.036)	
Lond. High	-	0.00	-	0.02^{*}
		(0.006)		(0.006)
Lond. Low	-	-0.03^{***}	-	-0.02^{**}
		(0.006)		(0.006)
Smog impact	0.02^{**}	-	-0.01	-
	(0.041)		(0.006)	
High impact	-	-0.03^{***}	-	-0.04^{***}
		(0.007)		(0.033)
Low impact	-	0.06^{***}	-	0.01^{*}
		(0.007)		(0.006)
N	6830	6830	6830	6830
R^2	0.00	0.00	0.00	0.00

Notes: 'Poor health' at fifty defined by lowest response to a three category question on current health. 'Poor health' at sixty defined by lowest two responses to a five category question on current health. 1952 cohort refers to those *in utero* during the smog. Estimated smog impact from the interaction of treatment area and membership of 1952 cohort. Standard errors clustered at county level: stars indicate significance at 5 and 1, and 0.1%.

As with other outcomes, there is a statistically significant but small difference in the London and England & Wales trends of around 0.1 of a percent. This is unlikely to change the interpretation of results because - for London and the 'high' polluted districts within London - the estimates are negative but not significantly different from zero. Results from the area affected by 'low' pollution, where foetal mortality appears to have been less significant, are too strong and precisely estimated to have been driven by this difference.

Table 1.7 shows results for employment. The effect on employment for all of London is -1% but is not precisely estimated. For the areas of London less affected by the smog, those observed at fifty are 4% less likely to be in employment. For the more affected areas of London, the estimate positive, but is not significantly different from zero.

Effects by gender

The surviving members of the cohort affected *in utero* by the smog were slightly (about 1%) more likely to be male than average. For those born in the areas most severely affected by the smog, the

Figure 1.16: Employment in 2001, by year of birth

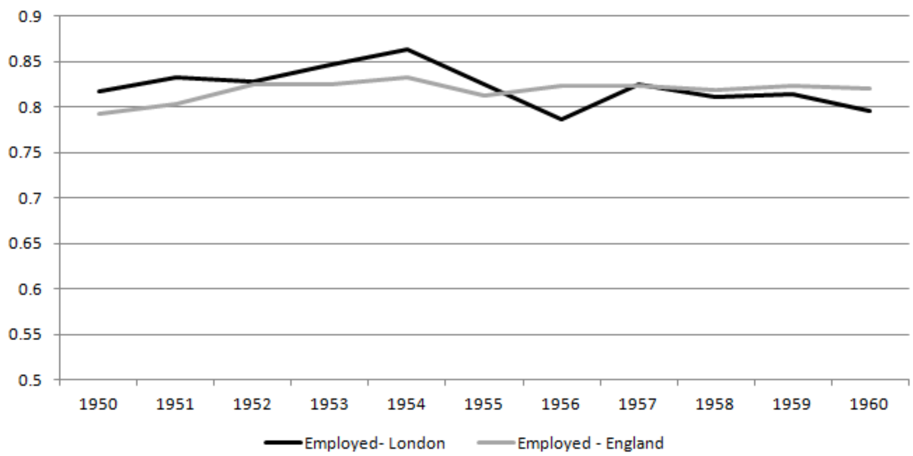

Notes: Source: Office of National Statistics Longitudinal Study

effect may have been quite strong - the 95% confidence interval for the estimate was [-1%, 10%]. To check whether results are being driven by a shift in the gender composition of the sample, the analysis was repeated for the two sub-samples. Overall, results were similar, but there were some differences in the effects on education and employment.

For education, the overall estimates were a 3% drop in the numbers with an A-level qualification. This result was driven entirely by changes in the academic achievement of girls, who were 7% less likely to have finished secondary school with this qualification. The effects for holding a degree were were the same for both genders, with those affected by the smog 5% less likely to hold a degree. For employment, the overall estimate was an imprecisely estimated drop of 1%. This was entirely driven by males in the sample, who were 4% less likely to be in employment than their peers, with the estimate significant at the 1% level.

1.5 Discussion

This paper aimed to study the short and long-term health impacts of foetal exposure the Great London Smog. The goal was to gather evidence about the health effects of high levels of pollution, such as those currently found in low- and middle-income countries.

Foetal Mortality

The first section used historical data on stillbirths to find whether exposure to the smog led to an increase in foetal mortality. Those *in utero* during the smog would have been born in 1953. In this year, the data showed a 2% increase in stillbirths in London, relative to national trends. As stillbirths often go unreported in official statistics, the exercise was repeated for live births. Results showed a 3% drop in live births against national trends, equivalent to 404 fewer births in London in the first nine months of 1953. Those surviving the smog were then observed in a 1% sample of the 1971 census (at around nineteen years old.) Expectations from Sanders and Stoecker (2011), and the papers cited within, were that foetal mortality would be more severe among boys - consistent with the Trivers-Willard Hypothesis that mothers might unknowingly favour female foetuses in hard times. For

Table 1.7: Employment: all London, and two levels within London

	Aged 50	
	London	Levels
1952 cohort	-0.01	-0.01
	(0.011)	(0.011)
London	0.04^{***}	-
	(0.009)	
Lond. High	-	0.04^{***}
		(0.008)
Lond. Low	-	0.05^{***}
		(0.008)
Smog impact	-0.01	-
	(0.011)	
High impact	-	0.02
		(0.010)
Low impact	-	-0.04^{***}
		(0.011)
N	6830	6830
R^2	0.00	0.00

Notes: 1952 cohort refers to those *in utero* during the smog. Estimated smog impact from the interaction of treatment area and membership of 1952 cohort. Standard errors clustered at county level: stars indicate significance at 5 and 1, and 0.1%.

London as a whole, there was no evidence that foetal mortality was more common for girls. In the districts of London that were most seriously affected by the smog, there was evidence that boys may have been *more* likely to survive than girls. People born in these area were 4% more likely to be male, with a 95% confidence interval of [-1%, 10%][23].

Information on deaths among those affected *in utero* by the smog showed them to be less likely to die than their peers. The major difference in the hazard rates between affected and unaffected cohorts was from age 45 onwards, suggesting that those not surviving *in utero* exposure would have been disproportionately likely to die in middle age. This is not, however, the only possible explanation for this pattern in the mortality data. It is also consistent with an explanation in which foetal mortality was not selective, but some foetal or maternal response to the smog improved later health - 'what doesn't kill you makes you stronger.'

[23]Although the results of this study are not supportive of the Trivers-Willard Hypothesis, they do not contradict existing studies finding an effect - much of the evidence on the Trivers-Willard hypothesis uses information from the U.S., where pollution levels are far lower than in London during the smog.

Outcomes for survivors

The first outcomes observed for survivors were educational attainment. Those affected were 3% less likely to hold an A-level, an effect driven almost entirely by a drop for females. Counter-intuitively, the estimated effect was weaker for the 'high polluted' districts of London. Both sexes from the affected cohorts were 5% less likely to hold a degree level qualification. Those affected by the smog were 2% more likely to report poor health at age 50. The figure for the 'low polluted' districts, in which foetal mortality appears to have been less severe, is 6%. Survivors exposed to the 'high polluted' areas were 3% *less* likely to report poor health than those in the control sample. Employment effects were similar, with those in 'low polluted' districts 4% less likely to be in employment, while there was some evidence that those *in utero* in 'high polluted' districts were more likely to be employed than peers.

What doesn't kill you

Taking information on the health, employment and educational prospects of survivors in to account, the hypothesis that a foetal or maternal response to the shock caused improved health seems less likely - those affected are, on average, in worse shape than their peers. The fact that negative health effects are 'stronger' in areas less affected by the smog also lends support to the strong-survivors hypothesis. It is plausible that stronger pollution produces a stronger effect on foetal mortality, and evidence on gender shifts in the 'high polluted' area lends support to this idea. If there was higher foetal mortality among those of poorer health, then survivors will be selected from particularly healthy individuals. Comparing these survivors to unaffected cohorts will result in a bias towards health.

Other episodes of severe pollution

Both of the areas of London labelled 'high polluted' and 'low polluted' in this paper experienced very high levels of atmospheric pollution. Levels were $800\mu g/m^3$ and $1800\mu g/m^3$ on average during the five days of the smog. Although pollution of this intensity essentially does not occur in high-income countries, there is evidence that it may occur quite frequently in low- and middle- income countries. Data from daily readings in Delhi[24], where the average level of pollution is currently around $280\mu g/m^3$, show how variable pollution levels can be. Between 2004 and 2010, there were fifty six occasions in which the Town Hall pollution meter recorded particulate levels over $800\mu g/m^3$, and nine occasions when the levels were over $1000\mu g/m^3$. The Town Hall meter provides observations for only around one day in ten and the true numbers are likely to be far higher. Neither Delhi nor India are exceptional in this respect - in 2012 over one hundred cities surveyed by the World Health Organisation had annual average pollution levels of over $100\mu g/m^3$. The number of people exposed is very large. One of the most polluted cities, Peshwar, is home to over three million people and has annual average pollution levels almost twice those of Delhi.

1.6 Conclusion

Two hundred years since the beginning of the industrial revolution, the health and economic consequences of polluting economic activity are still not fully understood. There is growing evidence that *in utero* exposure to atmospheric pollution can cause foetal mortality, and low birth weight for those

[24] Available from the Indian Central Pollution Control Board.

that survive. Much of the evidence on the foetal impact of atmospheric pollution comes from high-income countries where pollution levels are low, but where data on health and pollution are readily available. This paper studies the Great London Smog of 1952 in order to gain insights into the short and long-term consequences of foetal exposure to high levels of pollution. The essential approach of the paper is to compare the outcomes of people exposed to the smog *in utero* during 1952 (and born in 1953) to other cohorts in London, using unaffected counties of England & Wales to control for year-level effects. Evidence from historical registrar data showed that there were 2% more stillbirths in London in 1953, relative to national trends. As stillbirths are often unreported, this analysis was repeated with information on live births. Results showed a 3% reduction in the number of registered births in London, or around 400 fewer births in the first nine months of the year. Survivors were then identified by district and quarter of birth, and studied using the ONS-Longitudinal study, based on a 1% sample of individuals first observed in the 1971 census for England and Wales. In general, survivors observed fifty and sixty years after the smog were less likely to hold a formal qualification, less likely to be employed, and were generally in poorer health than their peers.

Bibliography

[1] Almond, D. (2006) fiIs the 1918 Influenza Pandemic Over? Long-Term Effects of In-Utero Influenza Exposure in the Post-1940 U.S. Population" *Journal of Political Economy* 114(4), 672-712

[2] Almond, D. & J. Currie (2011) fiKilling Me Softly: The Fetal Origins Hypothesis" *Journal of Economic Perspectives* 25(3), 153-172

[3] Arceo-Gomez, E. O., Hanna, R., & P. Oliva (2014) fiDoes the Effect of Pollution on Infant Mortality Differ Between Developing and Developed Countries? Evidence from Mexico City" *NBER Working Paper*, Number 18349

[4] Bell, M.L. & D.L. Davis (2001) fiReassessment of the Lethal London Fog of 1952: Novel Indicators of Acute and Chronic Indicators of Acute Exposure to Air Pollution" *Environmental Health Perspectives* 109(3)

[5] Bozzoli, C., Deaton, A., & C. Quintana-Domeque (2009) fiChild Mortality, Income, and Adult Height" *Demography* 46(4), 647-669

[6] Chay, K. Y. & M. Greenstone (2003) fiThe Impact of Air Pollution on Infant Mortality: Variation in Pollution Shocks Induced by a Recession" *Quarterly Journal of Economics* 118(3), 1121-1167

[7] Currie, J. & M. Neidell (2005) fiAir Pollution and Infant Health: What Can We Learn from California's Recent Experience? " *Quarterly Journal of Economics* 120(3), 1003-1030

[8] Currie, J. & R. Walker (2011) fiTraffic Congestion and Infant Health: Evidence from EZ-Pass" *American Economic Journal: Applied Economics* 3(1), 65-90

[9] Currie, J. & T. Vogl (2013) fiEarly-Life Health and Adult Circumstance in Developing Countries" *Annual Review of Economics* 5(1), 1-36

[10] Deschenes, O. (2012) fiTemprature, Human Health, and Adaptation: A Review of the Empirical Literature" *NBER Working Paper* Number 18345

[11] Deschenes, O. & M. Greenstone (2011) fiClimate Change, Mortality and Adaptation: Evidence from Annual Fluctuations in Weather in the U.S." *American Economic Journal: Applied Economics* 3(4), 152-185

[12] Isen, A., Rossin-Slater, M. & W. R. Walker (2014) fiEvery Breath You Take - Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970" *NBER Working Paper*, Number 19858

[13] Jan, J., Johansson, & J.P. Nilsson (2014) fiEconomic Status, Air Quality, and Child Health: Evidence from Inversion Episodes" *IZA Discussion Paper* Number 7929

[14] S. Jayachandran (2009) fiAir Quality and Early-Live Mortality Evidence from Indonesia's Wildfires " *The Journal of Human Resources* 44(4), 916-934

[15] Sanders, N. J. & C.F. Stoecker (2011) fiWhere Have All The Young Men Gone? Using Gender Ratios to Measure Fetal Death Rates" *NBER Working Paper* Number 17434

[16] Sanders, N. J. (2012) fiWhat Doesn't Kill You Makes You Weaker: Prenatal Pollution Exposure and Educational Outcomes" *Journal of Human Resources* 47(3), 826-850

[17] Stein, Z., Susser, M., Saenger, G. & F. Marolla (1975) *Famine and Human Development: The Dutch Hunger Winter of 1944-1945* New York: Oxford University Press

[18] Wilson / Ministry of Health (1954) fiMortality and Morbidity During the London Smog of December 1952" *Reports on Public Health and Medical Subjects No.95* London: Ministry of Health

[19] Zivin, G. J., & M. Neidell (2013) fiEnvironment, Health and Human Capital" *Journal of Economic Literature* 51(3), 689-730

Appendix

Health results using full categorical data

Table 1.8 shows differences-in-differences results from an an ordered logit model for self-reported categorical health. There are three categories for those aged fifty and five for those aged sixty, and so magnitudes are not directly comparable between ages. The ordered logit works on the odds of an individual choosing a particular category of health. The reported coefficients can be interpreted as the percentage change in the odds of being in a 'healther' category[25]. Results for London show that, compared to the general population, those affected *in utero* by the smog are 22% less likely, in expectation, to be in a 'healthier' category. The effect is in the same direction at sixty, but is not significantly different from zero. There are figures from the same individual viewed at different times - as the negative health effects presumably did not disappear after a decade, this might be explained by the comparison group 'catching up' in terms of poor health between fifty and sixty. Separating health effects for people affected by higher and lower levels of smog within London reveals large differences. In the area affected by weaker pollution, the effects are similar to those for the full sample - individuals are less healthy than their peers at both fifty and sixty. In the high pollution area that appears to have suffered the highest foetal mortality, survivors appear to be healthier than the general population at fifty, but are not significantly different at sixty.

[25] For example, if an individual had 8:1 odds of being in the highest health category, an estimated coefficient of -0.50 would imply a 50% drop in this odds ratio, to 4:1.

Table 1.8: Self reported health: all London, and two levels within London

	Aged 50		Aged 60	
	London	Levels	All London	Levels
1952 cohort	0.06	0.04	−0.02	−0.02
	(0.041)	(0.045)	(0.030)	(0.011)
London	0.42***	-	0.29***	-
	(0.041)		(0.036)	
Lond. High	-	−0.01	-	0.24***
		(0.049)		(0.039)
Lond. Low	-	0.60***	-	0.23***
		(0.049)		(0.039)
Smog impact	−0.22***	-	−0.02	-
	(0.041)		(0.030)	
High impact	-	0.23***	-	−0.01
		(0.048)		(0.033)
Low impact	-	−0.46***	-	−0.08*
		(0.043)		(0.033)
N	6830	6830	6830	6830
R^2	0.00	0.00	0.00	0.00

Notes: 1952 cohort refers to those *in utero* during the smog. Estimated smog impact from the interaction of treatment area and membership of 1952 cohort. Stars indicate significance at 5 and 1, and 0.1%. Self reported health in three categories for those observed in 2001 (at fifty) and in five categories for those observed in 2001 (at sixty.) Consequently, direction, but not magnitude of results are comparable between ages. Results from ordered logit can be interpreted as the percentage change in the expected odds of being in a 'healthier' health category.

Figure 1.17: UK wide emissions of smoke, millions of tonnes.

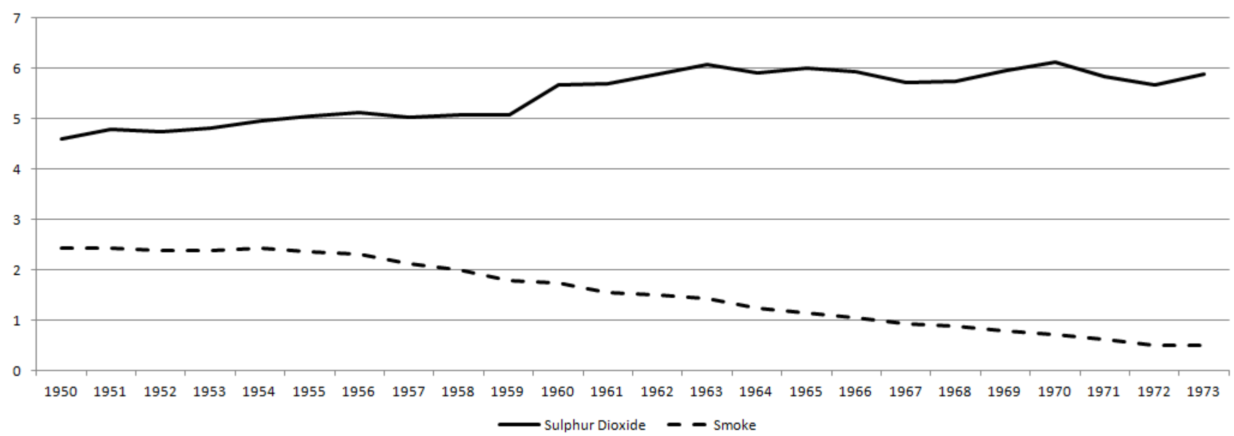

Notes: Data from the Fuel Research Board.

Figure 1.18: Daily minimum temperatures in London and England, degrees Celsius

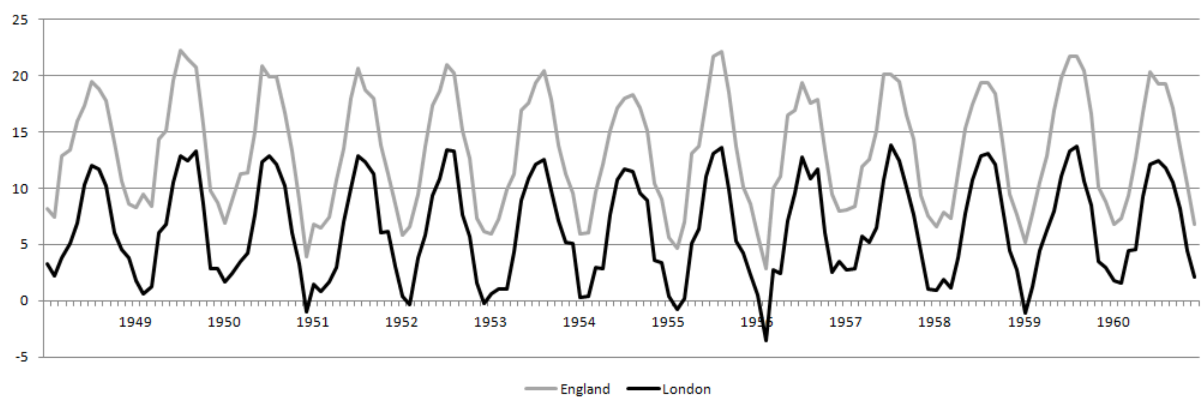

Notes: Temperatures measured in degrees Celsius. Data from the Met Office.

Figure 1.19: Labour market conditions (unemployment) in the UK

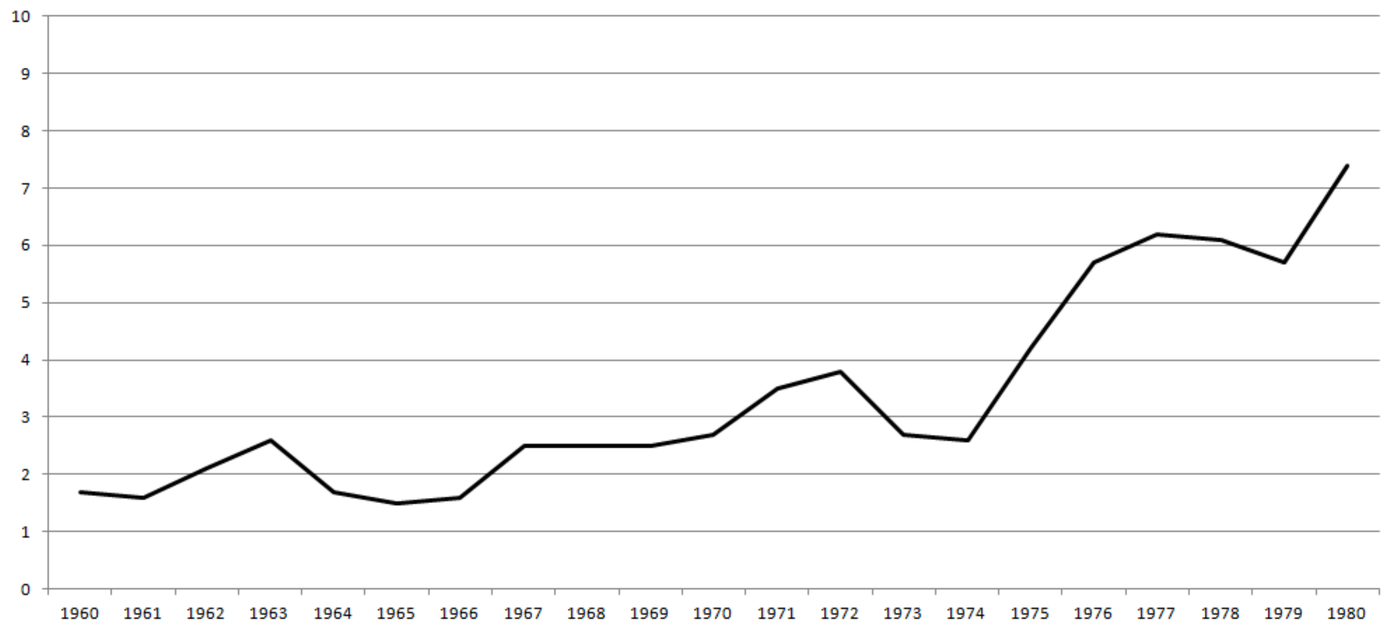

Notes: The (affected) 1952 cohort would enter the labour market between 1969 and 1971 when average levels of unemployment were 2.8% on average. The (unaffected) 1953 cohort would enter the labour market between 1970 and 1972 when average levels of unemployment were 3.1% on average. Source: Office of National Statistics

Chapter 2

Subject Choice and the Price of Education

2.1 Introduction

Although attending university is often an excellent investment, the outcomes for individual students can vary. A student graduating in 2011 with a degree in communications had a 11% chance of being unemployed one year after graduation. The equivalent figure for a student with a degree in nursing is 4%. For those students finding employment, the returns to a degree in the humanities are generally two percent lower than the returns to degrees in the STEM (science, technology, engineering, and maths) subjects[1]. This paper aims to find whether, as the private cost of attending university in England has risen, students have migrated to subjects, such as science or engineering, that are more highly valued in the job market.

Such a change would increase the financial return that students receive from university, but could also have wider effects. British businesses and public institutions such as the health service have long relied upon graduates from abroad to make up for skill shortages. For this reason, the promotion of eduction in the STEM subjects is a priority for the U.K. government. Were more costly education to increase the number of graduates with technical skills, the removal of educational subsidies could potentially benefit both the exchequer, and the wider economy.

An increase in the price of education might be expected to have two effects on the number of students studying STEM subjects. The price increase could result in a rise in demand for STEM subjects if students switched to technical subjects in search of a higher return. The price increase could result in a fall in demand if the higher fee discouraged participation of those already wishing to study a STEM subject. Data from the Higher Education Statistics Agency indicate that the 2006 fee increase had a significant effect on STEM subjects. Total enrolment, which had been growing at a rate of around five percent, dropped five percent in 2006, and didn't regain it's previous growth until 2008.

Using student-level data from the British Labour Force survey, this paper aims to find how much of this effect can be attributed to students switching courses. As the 2006 fee increase affected the participation of generally poorer, generally less academically able students, any causal estimate of the effect of fees on subject choice is likely to be affected by sample selection bias. Propensity matching is therefore used to obtain an estimate of the average treatment effect on the treated. Results indicate that the increase in fees caused a 5% decrease in the probability that a given student

[1] Walker et al (2011)

would study a STEM subject. Course level data indicate that this effect is concentrated in the *subjects allied to medicine* group, which includes degrees in pharmacology, nursing, and medical technology.

Overall, results indicate that students do respond to changes in the price of university education by switching to more lucrative degrees. However, as the most lucrative degrees are those in business and economics, the overall effect on STEM subjects was negative.

2.2 What affects subject choices at university?

From the perspective of the student, there are a two key differences between STEM subjects and other disciplines. The first is that STEM subjects are generally perceived to be more difficult. In the U.S., where students have greater latitude to change courses once enrolled, Stinebrickner and Stinebrickner (2013) conducted a longitudinal survey of students at Berea College. They found that students often arrived at university enthusiastic about studying scientific degrees, but changed courses due to the relative difficulty of the required courses. In the U.K., students in the STEM subjects do generally have more hours of classes than those in other subjects. However, the higher workload does not appear to have affected the probability that a student will fail to reach second year. As can be seen in figure 2.1, the drop out rates for STEM subjects vary greatly, but are not on average higher than those for subjects in the arts, or in the social sciences.

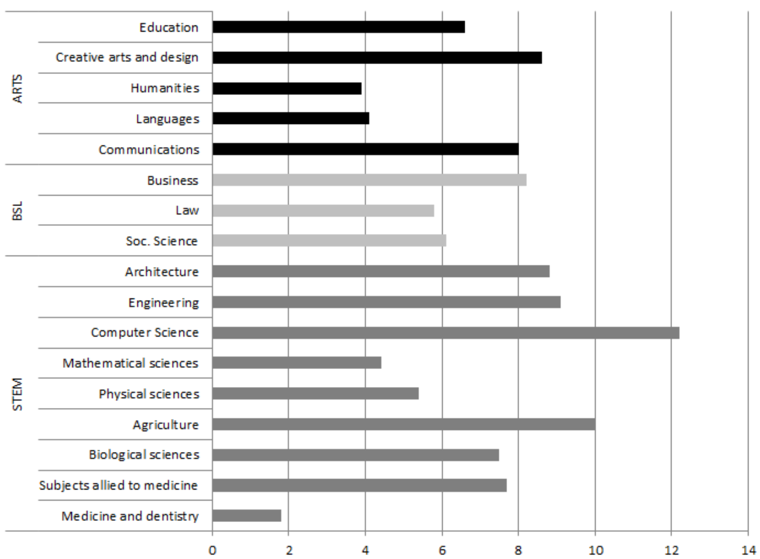

Figure 2.1: Percent of students not continuing to second year, 2010

The second difference is that graduates with degrees in STEM subjects are generally considered to perform better in the job market. In England, however, this is not always the case. Although the returns to degrees in the humanities are around 2% lower for all students, the returns to degrees in economics, management and law (LEM) are around the same as STEM degrees for women, and almost *twenty points* higher for men[2]. There are higher chances of employment in some subjects. As figure 2.2 shows, the chance of being unemployed is considerably lower for those studying topics relating to medicine.

[2]Walker et al (2011)

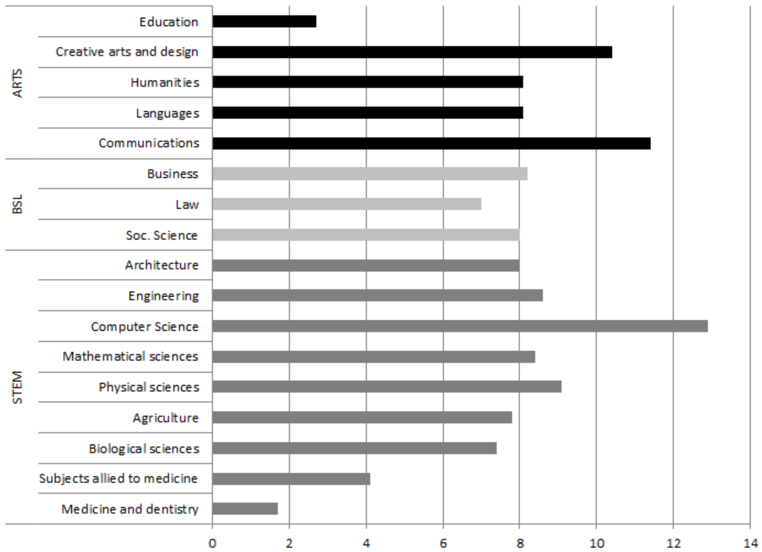

Figure 2.2: Unemployment one year from graduation, 2010

Evidence relating to how students actually choose their subjects is relatively sparse. Berger (1988) examined how future income streams affected major choice in American universities, finding that lifetime earnings were a more powerful predictor of outcomes than starting salaries. Montmarquette et al (2002), using data from the National Longitudinal Survey of Youth, modelled the choices of American students using both expected lifetime incomes and each student's probability of passing a particular course. They found the most important indicators of subject choice to be gender, with men much more likely to choose science degrees; and race, with non-white students generally choosing more lucrative degrees. The number of siblings that a student had increased the probability that they would pursue a degree in business or science, indicating that information transmission within the family might be influential. Lastly, they found that students that had taken loans tended to take degrees that were less technical, and consequently, less risky. Beffy et al (2012) studied the participation, subject choice, and education duration decisions of French students using data from the Gnration 92 and Gnration 98 surveys. Exploiting variation from the French business cycle, they found that although subject choice responds to changes in the relative wages of different career paths, the effect was very small.

The Effect of Fees

Why tuition fees might affect the subject choice decision can be demonstrated most easily with a simple model. Assume that, following Berger (1988), a student i's utility from subject j comes from a discounted stream of income Y^j minus the cost of education c, plus some idiosyncratic taste for that subject u_i^j. Letting V be a strictly concave function representing the student's utility from the monetary return to education, and letting $j \in \{\textsc{stem}, arts, work\}$, the student chooses the maximum of:

$$\begin{aligned} U_i^{\mathrm{STEM}} &= V(Y^{\mathrm{STEM}} - c) + u_i^{\mathrm{STEM}} \\ U_i^{arts} &= V(Y^{arts} - c) + u_i^{arts} \\ U_i^{work} &= V(Y^{work}) + u_i^{work} \end{aligned}$$

If we assume that $Y^{\text{STEM}} > Y^{arts} > Y^{work}$, a marginal student, for whom $U^{\text{STEM}} = U^{arts} > U^{work}$ would always respond to an increase in the price of education in one of two ways. First, a student would always switch to the more lucrative STEM subject. This is because, due to diminishing returns to income, the utility loss from fees has a larger effect on the utility from the less lucrative ARTS degree. Second, if tuition fees were sufficiently large that the utility from studying STEM dropped below the utility from working, students would exit from education altogether.

2.3 The English Education System

The majority of English students start school at six years old and receive around nine years of general education. At the age of 15, students specialise in up to eight GCSE subjects, which they study for a further two years. After receiving their GCSE results, students choose whether to stay at school, to work, to attend college, or to take an apprenticeship. Continuing students then generally choose up to four subjects to sit at AS-level, and up to three subjects to study in the following year for their A-levels[3].

Students apply for university while studying their A-levels with predicted grades. Applications are centralised though the Universities and Colleges Admissions Service (UCAS). Each application requires a small fee, the selection of up to five courses, and a personal statement. As the personal statement is usually focused on why the student is enthusiastic about their chosen subject, there is a strong incentive for students to apply for the same subject at five different universities. Relative to other higher education systems in Europe, English students are far more likely to apply for universities in other cities. Students then receive either unconditional offers, or offers conditional on their A-level results. Those not receiving an offer, or not meeting their conditions, are then free to approach universities though a second application stage known as *clearing*. Those exceeding their predicted grades are also allowed, at this stage, to apply to more prestigious institutions.

A first degree in England lasts three years. Institutions vary, but students are generally expected to stick to the subject that they applied to. Tuition fees and living expenses are paid though a mixture of loans, grants, and parental contributions. Grants are only available for the poorest students. Loans are government-backed, available to all students holding a place, and are almost interest free. Repayment is at a rate of 9% of any gross income above 15,000 and is generally organised though the tax system. The maximum loan available is reduced for students with wealthy parents, who are expected to cover any resulting shortfall in the student's finances.

Universities in England

University status in England has generally been granted by Royal Charter, a Papal Bull, or by an Act of Parliament. However, there have been two recent expansions of the system in 1992 and 2004, when two waves of colleges were permitted to take university status.

There are currently ninety one universities in England. All but two[4] are funded by the public Higher Education Funding Council for England (HEFCE) according to the following process. First, the government chooses the total amount to be allocated to higher education. The HEFCE then gives each university a single 'block grant' based on the number of students studying there, and the amount and

[3] The vast majority of students aiming for university study for A-levels. Some, generally fee-paying, schools prepare students for other qualifications such as the BTEC, the IB, or the Cambridge Pre-U.

[4] The exceptions are the University of Buckingham, which has charitable status, and the University of Law, which is a profit making institution.

quality of research generally produced[5]. The funding provided for each student is weighted by their subject taken, with lab-based teaching attracting around twice the funding as lecture-based teaching. Although universities are not expected to account for how they spend the block grant[6], those failing to attract the expected student numbers lose funding, and those that exceed their student quota are fined. Students also pay tuition fees, the level of which are set by the government. The incentives for universities are generally to maximise research output, to fill their allocated student places, and to attract the best candidates available.

Funding Changes in 2006

Tuition fees were first introduced in 1998, and were set at 1,000 a year. They were paid up front but, from 1999, could be added to the student's loan. The amount payable increased with inflation, reaching around 1,200 by 2005. From the 2006/7 academic year, universities were given the option of charging variable tuition fees to students, up to a maximum of 3,000. This upper limit turned out to be binding, with all universities choosing to charge the maximum allowed. As before, tuition fees were paid up front, but could be added to the student's loan so that - in principle - the increase in fees should not have prevented students attending due to liquidity constraints.

The effect of this policy change on enrolment can be seen in figure 2.3, which shows aggregate data from the Higher Education Statistics Agency (HESA) on the number of first year students enrolled in STEM and non-STEM subjects. Enrolment in STEM subjects, which had been growing at

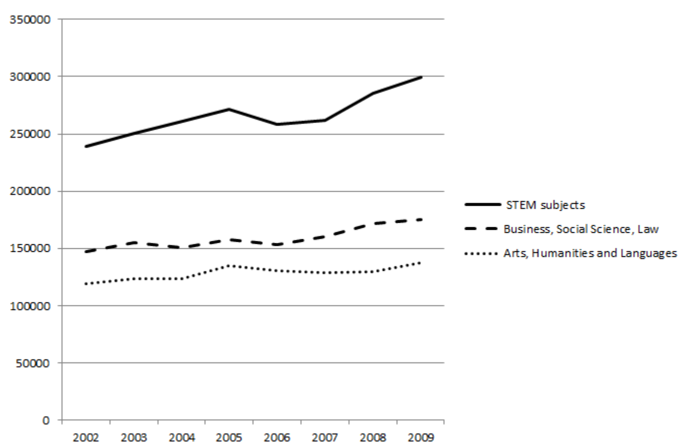

Figure 2.3: First year enrolment for STEM and non-STEM subjects

around five percent, shows a drop of five percent in 2006, only returning to trend in 2008. The policy change had a smaller effect on the other groups of students.

Of the STEM subjects, the most affected group was *subjects allied to medicine*, covering courses in anatomy, pharmacology, complementary medicine, nutrition, nursing and medical technology. As can be seen in figure 2.4, almost all of the variation in the STEM group is driven by changes in these subjects. Participation in all other subjects, including those not in the figure such as maths and vet-

[5]Research quality is determined by a research assessment exercise, which occurs about once every three years. Universities also receive funding from government research councils. The allocation of this funding is competitive, and the majority of funding goes to top institutions.

[6]The block grant accounts for the majority of university funding, but the HEFCE also allocates tied grants for infrastructure and programs aimed at extending participation.

erinary science, increased both in absolute and relative terms between 2002 and 2009. The exception is computer science, which experienced a drop in participation in both absolute and relative terms.

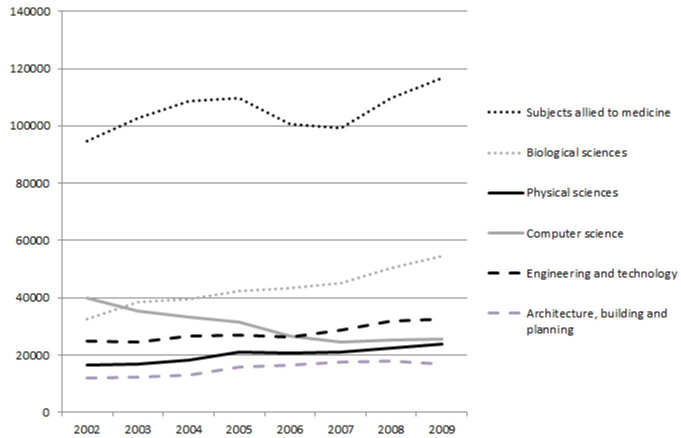

Figure 2.4: First year enrolment for selected STEM subjects

Participation in the non-STEM subjects was less affected by the funding changes. As can be seen in figure 2.5, art, history and philosophy, and languages all show a small bump in participation in 2005. This can be attributed to a change in deferral behaviour. The largest change was in business and administration, which showed flat or negative growth before 2006, and around three percent growth afterwards. The number of students training to be teachers rises steadily throughout the series, and the funding changes do not appear to have had an effect.

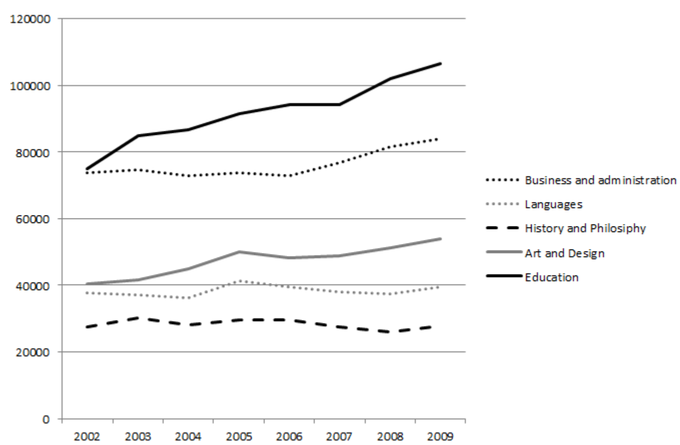

Figure 2.5: First year enrolment for selected non-STEM subjects

2.4 Data

Data comes from the 2003 to 2008 waves of the British Labour Force Survey. This is a large national survey in which around 60,000 households are interviewed for five consecutive quarters. Due to the available information on each student's family and household, it is possible to build a rich picture of their circumstances before enrolling in university. The sample consists of students aged seventeen and eighteen who have just enrolled in university. Those students for whom there is no parental information, and those that don't have A-level qualifications are removed from the sample. The final dataset contains 843 observations in total: 530 for the control group observed before 2006, and 313 for the treated group observed after 2005. Although there is information available about degree subject, subject choice will be divided into STEM and non-STEM for the main analysis. This is due mainly to the size of the treated and control groups.

2.5 Identification Strategy

This paper aims to find whether increasing the cost of education causes students to switch to STEM subjects. Variation in the cost of education comes from the 2006 increase in tuition fees from around 1,200 to 3,000 for all students attending university in England. Identification of the effect of tuition fees on subject choice requires that the *treated* students enrolling in university after 2006 and the *control* students studying before 2006 are comparable. There are two reasons why this might not be true. The first is the effect of tuition fees on university enrolment. Deardon et al (2011), using data from the British LFS from 1992 to 2007, estimate that an increase in tuition fees of 1,000 reduces university participation by 3.9%. They found significant factors affecting participation to be gender, ethnicity, parental education, parental income, and the region of England in which the student lives. It is therefore likely that the treated group contains a larger proportion of students, such as those with educated parents, that are less likely to be discouraged by tuition fees from participating in university. The second reason is deferred entry into university. Taking a gap year before university is relatively common in England, and it is likely that students getting their A-levels in 2005 would choose not to delay university in order to avoid the fee increase. Students taking gap years are generally from wealthier families, and one would therefore expect these students to be overrepresented in the sample of controls.

In the absence of a suitable natural experiment or instrument, a matching methodology will be used to manage these selection issues. The key assumption of this methodology is selection on observables. That is, that after conditioning on a set of observable variables, assignment to treatment is random. As the number of matching covariates will be large, an estimated propensity score will be used to match the treated with controls. This is a two-step procedure in which the probability of receiving treatment given observables is estimated, and then used to match the treated with controls that have a similar likelihood of treatment. Rosenbaum and Rubin (1983) demonstrate that if assignment to treatment given observables is unconfounded, so is assignment to treatment given the propensity score.

The CIA Assumption

As discussed above, the matching methodology relies critically on the assumption of selection on observables. Heckman et al (1997) stress the importance of using the same questionnaire relating to individuals in the same market when constructing matching estimators. The questions in the labour force survey do not vary during the sample period, and relate to the English education market. The variables used for matching in this study will be the same as those used by Deardon et al (2011) in their estimation of the effects of tuition fees on participation. As income is not available in the household version of the labour force survey, socio-economic status was used as a proxy. The relationship between these two variables can be seen below in figure 2.6. Following Montmarquette et al (2002), who found 'number of siblings' to be an important predictor of education decisions, a variable tracking family size was included. A second variable not included in Deardon et al (2011) is

Variable	Treated	Matched Controls	All Controls
Male	40%	46%	46%
Black	2%	1%	1%
Asian	9%	10%	10%
Family Size	1.8	1.6	1.6
Parent: Highest Income	21%	24%	26%
Parent: High Income	26%	29%	29%
Parent: Medium Income	12%	11%	11%
Parent: Degree	25%	25%	25%
Parent: HNC	10%	9%	9%
Parent: A-levels	7%	7%	7%
Parent: GCSE	22%	17%	17%
Parent: Arts Profession	1%	0%	0%
Parent: Business Profession	7%	7%	7%
Parent: Stem Profession	21%	24%	24%
Parent: Unemployed	3%	3%	3%
Region: North	25%	25%	25%
Region: Merseyside	3%	4%	4%
Region: East Midlands	9%	6%	6%
Region: West Midlands	11%	11%	13%
Region: Eastern	11%	11%	11%
Region: London	14%	15%	15%
Region: South East	19%	19%	18%
Region: South West	9%	8%	7%
Observations	313	530	547

Table 2.1: Characteristics of the treated and control groups

the occupation of the head of the family. This was included on the basis that children with parents in business would observe higher returns to education from their parents, while those with parents in the arts might be more likely to seek work experience in the creative industries. Four bands were used: creative industries, STEM industries, education, and business.

Table 2.1 below shows the characteristics of the treated sample, the control sample, and the control sample used in the analysis. In general, the control and matched samples are very similar. The largest difference is in gender, with males representing 40% of the treated sample and 46% of the control sample. To ensure that results are not being driven by this difference, the analysis will be repeated on the male and female subsamples.

Given the rich set of pre-treatment variables available, selection on observables does not seem unreasonable. However, it is impossible to explicitly test the conditional independence assumption. It is however possible to give some idea of how the results would be affected by a failure of the CIA. This sensitivity analysis shall be run using a methodology proposed by Ichino, Mealli and Nannicini (2008), which checks the robustness of the causal estimates to a specific failure of the CIA.

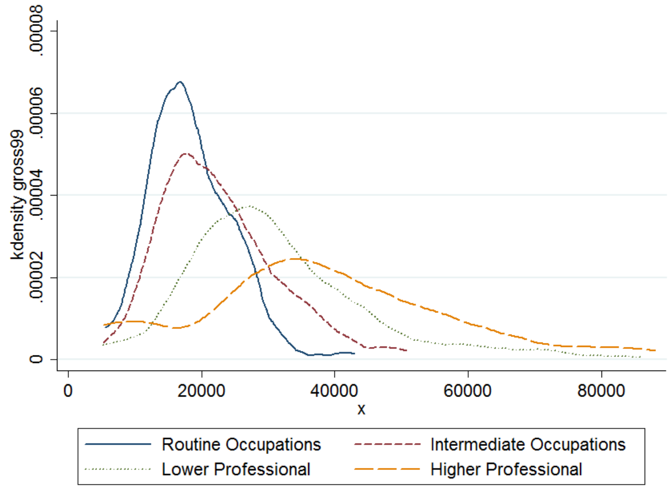

Figure 2.6: Socio-economic status and gross income, 40 to 60 year olds.

Exclusion Restrictions

A second key assumption is that nothing except for the tuition fee increase happened between 2003 and 2008 that might have affected the subject choice decision. Three possible challenges to this assumption are policy changes, changes in the job market, and changes in the university sector. These shall be discussed in turn.

Although there were no other changes to funding policy for English students, there were policy changes in other nations in the U.K. Tuition fees for Welsh students rose one year later than for English students and Welsh students studying in Wales were given a grant. In 2008, at the end of the sample, tuition fees for Scottish students were abolished. As the sample includes only English students, there will be no direct effect of these changes on English students. It is unlikely that any indirect effect would be large, as the proportions of Scottish and Welsh students at English universities, at around one and two percent, is relatively small.

Beffy et al (2012) demonstrate that the effects on subject choice of relative wage changes though the business cycle are very small. However, there were no large fluctuations in the period that might have influenced students. Youth unemployment started to tick up in early 2009, just outside of the period in question.

Beginning in 2001 with the University of Gloucestershire, the English university system underwent a large expansion when thirty one former colleges were given the right to rename themselves universities. Within the sample period, twenty new institutions gained university status. As almost all of these new universities had been collages beforehand, and already had degree awarding status, this policy change is unlikely to have affected subject choice. However, in order to control for any effect that these changes might have made, students with less than two A-levels, that are more likely to study in the new universities will be removed from the sample.

2.6 Estimation

Estimation of the Propensity Score

Estimation of the propensity score was conduced using the *pscore*[7] package. This program estimated the probability that an individual is assigned to treatment with a probit estimator. The propensity scores for the treated and control groups can be seen in figure 2.7 below[8]. As expected, the propensity distribution of the treated group is to the right of the distribution for controls. The region of common support is [0.16, 0.80]. For the estimation of the average treatment effect on the treated, it is important that there is a comparable control for each treated individual. This is the case, although the control distribution is a little sparse in the right tail. There are just seven observations with propensity scores above 0.7, only one of which is in the control group. As a robustness check, the main analysis was repeated with these observations removed, with no effect on results.

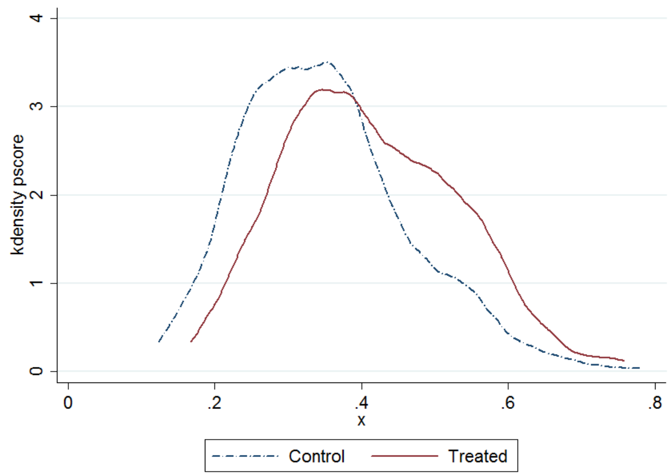

Figure 2.7: Propensity score for the treated and control samples

The *pscore* program also stratified the sample into blocks by their propensity scores and checked that each block is balanced. The final number of blocks chosen was six, containing 40, 220, 290, 270, 30, and 1 observations repetitively. All were balanced, though the final block contained no observations in the control group, and was dropped for estimation.

Choice of Matching Procedure

There are a number of alternative matching procedures available, offering different trade-offs between the number and quality of matches. This study uses four of the most popular: stratification, nearest neighbour matching, radius matching and kernel matching. Each is briefly discussed below[9].

Nearest neighbour matching works by matching each treated subject with the control with the closest propensity score. This study uses nearest neighbour matching with replacement, allowing controls to be used more than once. Advantages are that the quality of matching is potentially better

[7] Becker and Ichino (2002). This is a suite of programs available at the STATA webpage for estimating propensity scores, running diagnostics, and estimating by various modes of matching. See http://www.sobecker.de or http://www.iue.it/Personal/Ichino for more information.

[8] Regression results can be seen table 2.8 in the appendix.

[9] For a more detailed discussion, see Ichino and Becker (2002), or Caliendo and Kopeinig (2005)

than in other methods. The main disadvantages are that the nearest neighbour could be quite far away, and that some information from unused controls is lost. This method could work well with the sample in this study, as there exist good matches for most of the treated.

Radius matching solves the problem of poor matches by imposing a limit on the distance between the propensity scores of matched individuals. When there are multiple controls available for a treated observation, a weighted average is used. Advantages of this method are that poor matches are minimised, and that more of the sample is used than in nearest neighbour matching. Disadvantages are that matches could be worse on average than in nearest neighbour matching because all controls within the chosen radius are used to construct the counter-factual. As radius matching brings an improvement in the quality of matches over nearest neighbour matching only when treated observations without good controls are removed, the minimum radius shall be chosen so that it just binds. This will reduce the number of treated in the sample, but should improve the average quality of the match.

Matching by stratification works by dividing the sample into blocks by propensity score such that the balancing property is satisfied. As discussed earlier, the sample has been divided into six blocks, though the sixth block contains no controls and just one treated, and is dropped from the analysis. The advantage of stratification is that it uses almost all observations, even when the number of controls outnumber the treated. The disadvantage is that some blocks must be deleted in the case where suitable controls are not available, and that matches are generally of a lower quality than in radius or nearest neighbour matching.

Kernel matching matches each treated observation with a distance-weighted average of all controls. Advantages of this method are that it uses all observations in the sample. Disadvantages are that, as all observations are used, it is especially important that the treated and control observations share a common support. For this reason, all estimates shown in the following section are based on the subsample for which there is common support. As with stratification, the match will be on average poorer than in radius or nearest neighbour matching, but the variance of the estimated ATT should be lower.

Results

Results from the four matching estimators and an OLS estimate are shown in table 2.4 below. All estimates of the ATT are negative, suggesting that tuition fees resulted in students switching away from STEM subjects. Nearest neighbour matching, which should provide the most precise matching, indicates an effect that is considerably stronger than the other methods. The kernel matching, stratification and OLS estimates are broadly similar, indicating a five percent reduction in the probability of a student choosing a STEM subject. The estimate from radius matching lies between those of the other estimates. Course level data shown in figure 2.4 indicate that this effect is probably restricted entirely to subjects in the *subjects allied to medicine* category, including nursing, pharmacology, and medical technology. Figure 2.5 indicates that these students might be transferring to courses in business or education.

2.7 Sensitivity Analysis

The effect of gender

In the discussion of the identification strategy, it was noted that the percentage of men in the treated sample was 6% larger than in the control. In order to see if this is driving results, a sub-sample

Estimator	Treated	Control	ATT	p
Nearest Neighbour Matching	313	272	-0.11^{*}	0.02
			(0.046)	
Radius Matching	260	399	-0.07	0.07
			(0.039)	
Stratification	312	531	-0.05	0.12
			(0.034)	
Kernel Matching	313	530	-0.05	0.12
			($0.040^{\dagger}$)	
OLS	313	531	-0.05	0.11
			(0.046)	

Table 2.2: The effect of tuition fees on the probability of choosing a STEM subject.
† indicates bootstrapped standard errors. Matching radius: 0.001. Detailed OLS results can be seen in table 2.7 in the appendix.

analysis was conducted. Results can be seen in the table below.

Estimator	Treated	Control	ATT	p
M- Nearest Neighbour Matching	125	119	-0.18^{*}	0.01
			(0.073)	
F- Nearest Neighbour Matching	188	159	-0.08	0.17
			(0.056)	
M - Radius Matching	85	149	-0.10	0.14
			(0.069)	
F - Radius Matching	143	181	-0.03	0.56
			(0.054)	
M - Stratification	125	245	-0.06	0.26
			(0.055)	
F - Stratification	187	286	-0.05	0.23
			(0.042)	
M - Kernel Matching	125	245	-0.04	0.46
			($0.053^{\dagger}$)	
F - Kernel Matching	188	285	-0.05	0.23
			($0.044^{\dagger}$)	
M - OLS	125	245	-0.05	0.34
			(0.053)	
F - OLS	188	285	-0.05	0.25
			(0.042)	

Table 2.3: The effect of tuition fees on the probability of choosing a STEM subject, by gender. † indicates bootstrapped standard errors. Matching radius: 0.05

With both nearest neighbour and radius matching, the negative effect for men was larger than for women. However, for both genders the estimated ATTs were of the same sign and similar magnitude to those in the main results. Although there is evidence that there are differences between the genders, they are probably not driving the main results.

Estimating on 'thick' support

Figure 2.7 shows that the right hand tail of the common support is thin, containing few controls. Following the suggestion of Black and Smith (2004), the analysis will be re-run with just the 'thick' support, in order to see if results were being driven by the right hand tail. Specifically, all observations with estimated propensity scores outwith the closed inverval [0.20, 0.55] were dropped from the sample. As can be seen in the table below, estimating in only the area with 'thick' support made little difference to the estimated coefficients, or to their standard errors.

Estimator	Treated	Control	ATT	p
Nearest Neighbour Matching	265	249	-0.10^{*}	0.03
			(0.048)	
Radius Matching	234	373	-0.07	0.10
			(0.041)	
Stratification	265	463	-0.05	0.14
			(0.036)	
Kernel Matching	265	463	-0.04	0.21
			($0.034^{\dagger}$)	
OLS	265	463	-0.06	0.11
			(0.035)	

Table 2.4: The effect of tuition fees on the probability of choosing a STEM subject. Estimating on 'thick' support. † indicates bootstrapped standard errors. Matching radius: 0.001

Testing the sensitivity of the estimated ATT to unobserved confounders

Although it is impossible to test the conditional independence assumption, it is possible to give some indication of how sensitive the estimated effects are to an unobserved confounder. Ichino et al (2008) propose the following methodology. First, assume that the selection on observables assumption does not hold because of an unobserved binary variable U that is missing from the set of controls W. Formally, letting T denote treatment status, and letting Y_i denote the (binary) outcome given treatment status,

$$Pr(T = 1|Y_0, Y_1, W) \neq Pr(T = 1|W)$$
$$Pr(T = 1|Y_0, Y_1, W, U) = Pr(T = 1|W, U)$$

Second, the binary confounding factor is characterised by four parameters p_{ij} that specify the probability that $U = 1$ for an individual with outcome j and treatment status i.

$$Pr(U = 1|T = i, Y = j, W) = Pr(U = 1|T = i, Y = j) \equiv p_{ij}$$

The distribution of U conditional on T and Y is assumed not to vary with W. Third, given the values of p_{ij}, each subject is assigned a value of U given their treatment i and outcome j status. It is now possible to re-estimate each subject's propensity score, and to estimate the ATT. This procedure is then repeated a large number of times with the same set of p_{ij}s. The final ATT given by the procedure is an average of the ATTs implied by each draw. The major choice in this procedure are the values of p_{ij}. In this exercise, five sets of values are used. The first set is that of a neutral confounder, the other four are chosen to mimic the behaviour of known variables. The analysis will be run twice.

	p_{11}	p_{10}	p_{01}	p_{00}	Outcome effect	Selection effect	ATT	s.e.
No confounder	0.00	0.00	0.00	0.00	-	-	-0.11	0.046
Neutral confounder	0.50	0.50	0.50	0.50	1.015	1.016	-0.08	0.054
Confounder like								
Male	0.51	0.36	0.55	0.42	1.779	0.766	-0.06	0.054
Asian	0.14	0.07	0.12	0.09	1.605	0.899	-0.09	0.053
High Income	0.17	0.22	0.23	0.27	0.838	0.767	-0.08	0.054
Parent has GCSE	0.07	0.07	0.07	0.06	1.304	1.170	-0.09	0.052
London Region	0.14	0.14	0.12	0.16	0.749	0.967	-0.09	0.053

Table 2.5: Sensitivity of nearest neighbour estimates to an unobserved binary confounder

Once for nearest neighbour matching, which showed a stronger effect than the other methods, and once for kernel matching, which agreed best with the results from the other methods.

For concreteness, let the posited unobserved variable U be a dummy taking the value one if the subject is hard working. As discussed above, the four values of p in table 2.5 show the probability that an individual is hard working. For example, p_{11} shows the probability that an individual who goes to university after 2006 and studies a STEM subject is hard working. The outcome effect is the effect that being hard working has on the untreated outcome, all other things being equal. In this case, this is the effect that being hard working has on the probability of studying a scientific degree before 2006. The selection effect shows the effect of being hard working on assignment to treatment, all other things being equal. In this case, it shows the effect that being hard working has on the probability that someone is in the post-2006 cohort. Results are reported for a neutral confounder, and then for confounders with characteristics similar to the male, asian, high income, gcse, and london region dummies. The specification of U that had the largest effect on the ATT was the one mimicking the distribution of the male dummy. This reduced the ATT by five percentage points from -11% to -6%. In general, all specifications of the confounder reduced the magnitude of the ATT, but estimates remained clearly negative, and within the bounds of the five main estimated ATTs.

The same exercise was repeated with the kernel matching estimator. This estimator was chosen for the sensitivity analysis because its central estimates had the smallest magnitude, and were closer than those from the nearest matching estimator to the results indicated by the stratification and OLS estimators. As before, the specification of U that had the largest effect on the ATT was the

	p_{11}	p_{10}	p_{01}	p_{00}	Outcome effect	Selection effect	ATT
No confounder	0.00	0.00	0.00	0.00	-	-	-0.05
Neutral confounder	0.50	0.50	0.50	0.50	1.026	1.015	-0.05
Confounder like							
Male	0.51	0.36	0.55	0.42	1.793	0.768	-0.04
Asian	0.14	0.07	0.12	0.09	1.590	0.917	-0.05
High Income	0.17	0.22	0.23	0.27	0.835	0.772	-0.05
Parent has GCSE	0.07	0.07	0.07	0.06	1.289	1.191	-0.05
London Region	0.14	0.14	0.12	0.16	0.774	0.951	-0.05

Table 2.6: Sensitivity of kernel matching estimates to an unobserved binary confounder

one mimicking the distribution of the Male dummy. This reduced the ATT by one percentage point from −5% to −4%. The results from the kernel matching estimates showed very little variation to the other confounders, with all estimates remaining almost identical in sign and magnitude to the original result.

2.8 Discussion

A naive OLS estimate of the effect of tuition fees on the percentage of students studying STEM subjects shows a drop of around 5%. The main results of this study show that the tuition fees made individual students around 7% less likely to study STEM subjects. Taken together, these results indicate that the 2006 increase in tuition fees from 1,200 to 3,000 had two effects on subject choice. The first was that students switched away from STEM subjects. The second is that the effect of tuition fees on participation was probably a little stronger for students intending to study non-STEM subjects.

As discussed earlier, course level participation data from the HESA indicate that most of this effect was driven by the *subjects allied to medicine*. This group includes courses in anatomy, pharmacology, complementary medicine, nutrition, nursing and medical technology. The only other subject group that seemed to be seriously affected by the 2006 increase in tuition fees was *business and administration*, which moved from flat or negative growth to positive growth. Participation over time for the two subjects is shown in figure 2.8. Although the trend lines drawn are simple extrapolations,

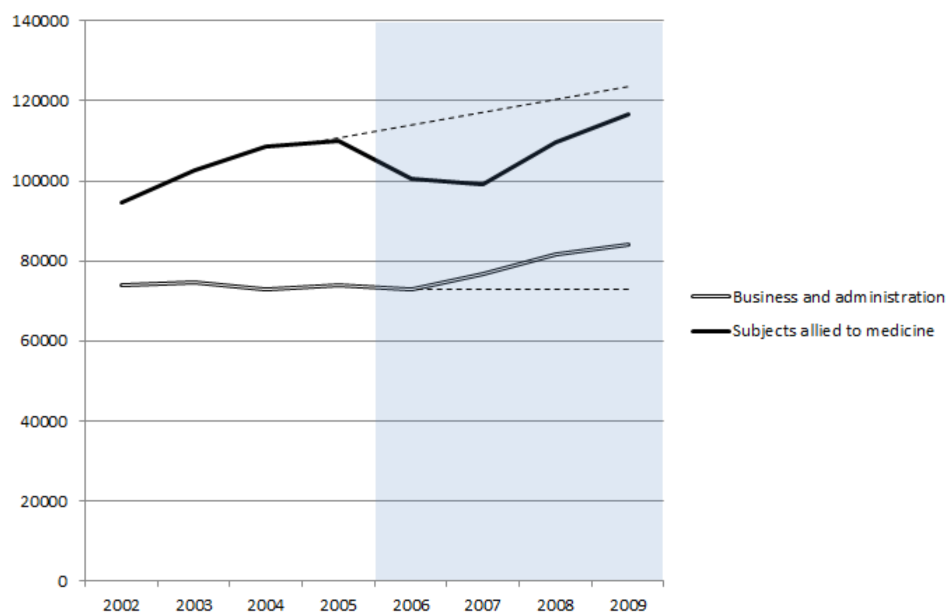

Figure 2.8: Participation over time

it does seem plausible that the increase in tuition fees resulted in a switch from subjects such as nursing, to subjects such as business. One explanation for why students would make this switch is the rate of return to the two subjects. Walker (2011) estimates that the return to degrees in business, social sciences, and law are similar to those for STEM degrees for girls, but around twenty points higher for boys. This explanation cannot explain all of the switching behaviour however. Although the nearest neighbour estimates of the switching behaviour for boys showed a stronger response than for girls, the other estimators showed similar responses. A second explanation could be that, as courses became more expensive, students avoided technical courses because they were afraid of failing. Although the actual probabilities of dropping out for business and subjects allied to medicine are almost identical, it is probably true that scientific subjects are still perceived to be more difficult.

The welfare impact of these changes should positive for at least the male students, who would be expected to earn a better return from a business degree than from a scientific degree. However, if the government wished to increase the supply of nurses and medical technicians in the economy, these results are discouraging. Possible remedial policies could be to more heavily subsidise the tuition fees of students studying nursing or similar subjects, or to increase the wages earned in these professions in the health service, increasing the return to this kind of degree.

2.9 Conclusion

Using student-level data from the British LFS, this study aimed to find whether a higher price for university education caused students to study subjects like science, technology, engineering or mathematics, that are more highly valued in the job market. Variation in the price of education came from the 2006 increase in tuition fees in English universities from around 1,200 to 3,000. A matching methodology was used in order to separate out the effects of this policy change on participation from those on subject choice. Results indicated that the increase in fees caused students to switch away from scientific subjects. Estimates from the matching estimators indicated that students paying fees were 5% less likely to study a scientific subject. Almost all of this effect was limited to subjects such as nursing, anatomy, pharmacology and medical technology and the most plausible candidate for the subject that students switched to was business and administration. Two explanations for this change present themselves. The first is that, although the returns to scientific subjects are generally higher than those in the humanities, the returns for business degrees are higher still. The second is that, although drop out rates are very similar for all subjects, scientific subjects are perceived to be more difficult. Students afraid of failing a course with student debt might be discouraged by higher fees from taking a chance.

Bibliography

[1] Beffy, M., Fougre, D., and Maurel, A., (2012) fiChoosing the field of study in post-secondary education: do expected earnings matter?” *The Review of Economics and Statistics* 94(1), 334-347

[2] Becker, S.O., Ichino, A., fiEstimation of average treatment effects based on propensity scores” *The Stata Journal* 2(4) 358-377

[3] Berger, M. C., (1988) fiPredicted future earnings and choice of college major” *Industrial and Labor Relations Review* 41,418-429

[4] Black, D., Smith, J., (2004) fiHow robust is the evidence on the effects of college quality? Evidence from matching” *Journal of Econometrics* 121, 99-124

[5] Caliendo, M., and Kopeinig, S., (2004) fiSome practical guidance for the implementation of propensity score matching” *DIW Berlin Discussion Paper 485*

[6] Heckman, JJ., Ichimura, H., Todd, P., (1997) fiMatching as an econometric evaluation estimator: evidence from evaluating a job training programme” *Review of Economic Studies* 64, 605-654

[7] Ichino, A., Mealli, F., and Nannicini, T., (2008) fiFrom temporary help jobs to permanent employment: what can we learn from matching estimators and their sensitivity?” *Journal of Applied Econometrics* 23, 305-327

[8] Lorraine Dearden, L., Fitzsimons, E., and Wyness, G., fiThe Impact of Tuition Fees and Support on University Participation in the UK”, *Centre for Economics of Education Departmental Paper 126* Available at www.cee.lse.ac.uk

[9] Montmarquette, C., Cannings, K., and Mahseredjian, S., fiHow do young people choose college majors?” *Economics of Education Review* 21, 543-556

[10] Rosenbaum, P.R. and D.B. Rubin, fiThe Central Role of the Propensity Score in Observational Studies for Causal Effectsfi, *Biometrika*, 70 41-55

[11] Walker, I., and Zhu, Y, (2011), fiDifferences by degree: Evidence of the net financial rates of return to undergraduate study for England and Wales, *Economics of Education Review*, 30, 1177-1186

[12] Stinebrickner, R., and Stinebrickner, T.,(2013), fiA Major in science? Inital beliefs and final outcomes for college major and dropout”, *NBER Working Paper 19165*

2.1 Appendix

Table 2.7: Results from OLS regression. N = 860, R^2 = 0.06

Variable	Coefficient	p
Treatment	-0.0522	0.114
	0.032994	
Male	0.107024	0.001
	0.031843	
Black	-0.12928	0.312
	0.127752	
Asian	0.127688	0.039
	0.061862	
Parent: STEM occupation	0.030558	0.47
	0.042297	
Parent: Medium Income	-0.0228	0.697
	0.05857	
Parent: High Income	-0.04781	0.328
	0.048848	
Parent: Highest Income	-0.10765	0.041
	0.052709	
Parent: Unemployed	-0.08262	0.382
	0.094499	
Parent: GCSEs	-0.03303	0.486
	0.047414	
Parent: A-levels	0.064583	0.342
	0.06796	
Parent: HND	0.121614	0.05
	0.061954	
Parent: Degree	0.096482	0.053
	0.049774	
Family Size	0.049707	0.01
	0.019129	
Region: North	-0.07896	0.236
	0.066602	
Region: Merseyside	-0.04532	0.656
	0.101846	
Region: Midlands	-0.03272	0.661
	0.074669	
Region: Eastern	0.013238	0.861
	0.075531	
Region: London	-0.09502	0.197
	0.073622	
Region: Southeast	-0.04186	0.545
	0.069049	
Region: Southwest	0.031434	0.698
	0.081101	
Constant	0.226963	0.003
	0.075375	

Table 2.8: Propensity Score Estimation: N = 860; R^2 = 0.05

Variable	Coefficient
Male	-0.16587
	0.091562
Black	0.006945
	0.356188
Asian	-0.09491
	0.178017
Parent: Medium Income	-0.23719
	0.168539
Parent: High Income	-0.4573^{***}
	0.14346
Parent: Highest Income	-0.59911^{***}
	0.16046
Parent: Creative	1.184217
	0.641565
Parent: Business	0.233147
	0.193918
Parent: Education	-0.3417
	0.237026
Parent: STEM	0.053992
	0.143345
Parent: Unemployed	-0.31767
	0.272085
Parent: GCSEs	0.416518^{*}
	0.134376
Parent: Alevels	0.401628^{*}
	0.19494
Parent: HND	0.417318^{*}
	0.178048
Parent: Degree	0.4554^{*}
	0.144122
Family Size	0.225388^{*}
	0.053875
North	-0.24586
	0.188027
Merseyside	-0.34981
	0.292967
West Midlands	-0.30876
	0.210833
Eastern	-0.23976
	0.21338
London	-0.27607
	0.209102
South East	-0.21367
	0.195002
South West	-0.12349
	0.230369
_cons	-0.39132
	0.209131