Empirical Essays in Family Economics

Submitted by

Tanya Wilson

for the degree of Doctor of Philosophy

of the

Royal Holloway, University of London

2015

Declaration

I, Tanya Wilson, hereby declare that Part I of this thesis, comprising Chapters 1, 2, and 3 are entirely my own research work. Part II is comprised of collaborative research. Chapter 4 was conducted in collaboration with Dr. Timo Hener. My contribution to this paper included writing the program code used for the estimation, analysing the data and writing the paper. Chapter 5 is co-authored work with Professors Dan Anderberg, Helmut Rainer and Jonathan Wadsworth, where my contribution entailed participating in discussions to exchange research ideas, preparing the data for analysis and estimation.

Signed .... (Tanya Wilson)
Date:

Forwards

Abstract

This thesis addresses questions in the field of Family Economics regarding the effect of education on teenage motherhood and choice of spouse, and the influence of local labour market conditions on domestic violence.

The first part of the thesis presents evidence regarding the effect of education policy on teenage fertility. Chapter 1 analyses the impact of an exogenous increase to education, induced by a legislative change to the minimal schooling requirement, on the propensity for early motherhood. The findings indicate strong evidence that the schooling reform had a substantial downwards impact on adolescent fertility, which persisted beyond the new compulsory school leaving age. The second chapter presents an evaluation of the Education Maintenance Allowance, a conditional cash transfer program implemented to increase the post-compulsory education participation rates of young people from disadvantaged families. The results reveal that the programme was successful at increasing participation rates, with a stronger response observed for teenage boys. This increase in participation had a significant impact on teenage motherhood, driven equally by a decline in the number of conceptions and an increase in the abortion rate. Weaker, but suggestive evidence is found that the programme mitigated youth crime.

The second part of the thesis addresses marriage and partnership questions. The third chapter investigates the impact of an education reform which induced a cohort discontinuity in the level of qualifications received by individuals, and finds that cohorts in the neighbourhood of the reform threshold cannot achieve typical matching patterns. Specifically, spouses choose smaller age gaps and accept differently qualified partners. The fourth and final chapter examines how changes in unemployment affect the incidence of domestic abuse. Combining data on individual experience of intimate partner violence with locally disaggregated labour market data, the analysis shows that a womans propensity to experience partner abuse decreases with male unemployment, but increases with female unemployment.

Acknowledgements

I would like to express my gratitude to a number of people at Royal Holloway who have encouraged and supported me in completing this thesis. First and foremost I would like to thank my supervisors, Dan Anderberg and Arnaud Chevalier, for all their advice and guidance. I am also deeply grateful to Jesper Bagger, Dan Hamermesh, Melanie Lührmann, Andrew Mountford, Philip Neary, Juan Pablo Rud, Andrew Seltzer and Ija Trapeznikova for helpful comments and kind words of encouragement.

To Dan Anderberg, Helmut Rainer, Timo Hener and Jonathan Wadsworth for allowing me to include our papers in this thesis.

To Jo Hible, Jeanne Johnson, Louise Goodwin, Nirusha Vigi and Nikki Webb who all indirectly helped me with this dissertation.

To all my fellow PhD students in London, and especially Chiara, Claudia, Cormac, Francisco, Lena, Michael, Myra, Nic, Sefi, Simon, Theo, Warn and Vusal, who made the PhD process enjoyable.

To the rugby girls who keep me grounded: Alex, Clanky, Jo Bev, Michelle, Sinners, Strudders and Viks.

To the Burwood clan, Emmajane, Richard, Chlöe, Grace, Maximus and Sampson, who are my second family.

And finally, to my parents for their enduring faith in me.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis addresses questions in the field of Family Economics in two parts, each comprised of two chapters. In Part I evidence is presented regarding the effect of two different education policies on teenage fertility. First the influence of an educational reform which increased the duration of compulsory schooling is considered, secondly a programme to increase post-compulsory participation in education is analysed. Part II of the thesis addresses marriage and partnership questions, specifically investigating marriage market responses to a cohort-level education reform, and the impact of local labour market conditions on the incidence of intimate partner violence.

The first chapter asks whether education policy can provide an effective tool to reduce the incidence of teenage motherhood. Data from the largest UK household-level survey is used to investigate the impact of a change in legislation implemented by the UK government in the early 1970s known as the Raising of School Leaving Age (RoSLA), which increased the duration of compulsory schooling by one year, on the timing of fertility using a regression discontinuity design. The findings indicate strong evidence that the schooling reform induced a downwards impact on fertility not only at the age at which the reform 'bites', but also persisted beyond the new compulsory school leaving age. Overall the analysis suggests that the increase in mandatory education caused a postponement of fertility over the teenage years only with the influence of the reform dissipating after age 20.

A different type of education policy is examined in Chapter 2. Specifically the impact of a conditional cash transfer programme, the Education Maintenance Allowance, which also incorporated elements of a financial incentives programme is examined. This programme was implemented in the United Kingdom in order to encourage participation in post-compulsory education by young people from low-income households. The results indicate a significant positive impact of the programme on staying-on rates, with a stronger response observed for teenage boys and those individuals eligible for the maximum allowance, who have the poorest familial pecuniary conditions. This increase in participation induced a substantial reduction in teen maternity rates, driven almost equally by a decline in conceptions and an increase in the abortion rate. Weaker, but suggestive evidence is found that the programme influenced youth crime.

The third chapter investigates another implication of the RoSLA legislative change to the duration of compulsory education, the reform also considered in Chapter 1, which induced a cohort discontinuity in the level of qualifications received by individuals. As marital matches around the world are characterised by two pervasive features: first women marry older men, second spouses have similar education levels, this cohort discontinuity implies that typical matching patterns cannot be achieved for individuals born close to the reform implementation date. Specifically, spouses choose smaller age gaps and accept differently qualified partners. As their potential partners are from untreated cohorts, women in early reform cohorts form marriages with lower positive assortative matching on education, whereas men can increase the degree of homogamy as their likely spouses are also subject to the reform. These opposing effects imply gender heterogeneity in education reform effects.

The fourth and final chapter examines how changes in unemployment affect the incidence of domestic abuse. A theory is developed which predicts that male and female unemployment have opposite-signed effects on domestic abuse: an increase in male unemployment *decreases* the incidence of intimate partner violence, while an increase in female unemployment *increases* domestic abuse. Combining data on individual experience of intimate partner violence with locally disaggregated labour market data, the analysis provides strong evidence in support of the theoretical prediction, and shows that a womans propensity to experience partner abuse decreases with male unemployment, but increases with female unemployment. This latter chapter is forthcoming in the *Economic Journal.*

Part I

Education Policy and Teenage Fertility

Chapter 2

Compulsory Education and Teenage Motherhood

2.1 Introduction

Teenage motherhood is widely regarded as an important socio-economic issue for two key reasons. First, individuals who are restricted in their human capital investment in adolescence may not reach their lifetime's economic potential. Second, there is an important inter-generational dimension associated with early childbearing, as the children born to teenage mothers tend to have poorer outcomes, and also themselves have a higher probability of becoming mothers at an early age (Paniagua and Walker, 2012).

These direct and indirect consequences of teenage motherhood have been widely documented. Analysis based on observed differences between women who gave birth as a teenager and women who became mothers at an older age find substantial adverse effects of early childbearing on a number of lifetime outcomes, such as lower levels of educational attainment (Moore and Waite, 1978; Klepinger, Lundberg, and Plotnick, 1995), and significantly higher rates of poverty, welfare receipt and lower household income (Bronars and Grogger, 1994; Ermisch and Pevalin, 2003). Studies that address the potential endogeneity of the fertility decision to ascribe a causal effect of adolescent motheerhood reveal somewhat disparate results both with respect to the impact on the mother herself (see, *inter alia*, Chevalier and Viitanen, 2003; Hotz, McElroy, and Sanders, 2005; Fletcher and Wolfe, 2009; Ashcraft, Fernández-Val, and Lang, 2013) and the outcomes of the child (Francesconi, 2008). Such analyses with arguably a closer comparison group tend to indicate that teen mothers would have poorer economic outcomes even if they had delayed motherhood. Indeed, Kearney and Levine (2012) suggests that for some individuals the decision to become a young mother is a rational choice in response to low expectations of future economic opportunities, rather than an unintentional consequence. In short, the weight of evidence overwhelmingly points to the existence of negative effects both in the short and long-run. Hence interventions that mitigate adolescent fertility rates are regarded as plausible mechanisms through which to improve the life trajectories of young women whom project a high proclivity toward teenage motherhood.

Given the interdependence of the education and early fertility decision, one potential channel of influence is education policy. In the context of mitigating teenage motherhood interest lies in the ability of the institutional environment to affect the timing of fertility. Exogenous differences in the duration of mandatory schooling have been used to elicit the causal effect of education on the likelihood of becoming a teenage mother. Changes in legislation regarding the minimum age at which an individual becomes eligible to leave school were first used by Black, Devereux, and Salvanes (2008) to investigate the effect of education on teenage fertility[1]. The authors propose two mechanisms through which the legislation changes exert an effect on fertility. First the "incarceration effect", which in the spirit of the findings of Jacob and Lefgren (2003) regarding the impact of schooling on youth crime, can be understood that as individuals are required to remain at school for one year longer, this reduces the opportunity to engage in risky activities, which leads to downward pressure on their fertility. Second, the "human capital effect", whereby individuals reduce their fertility in response to receiving more education and hence better labour market prospects as a result of the legislation change.

Using data for both the US and Norway - two countries with very different institutional environments - they obtain remarkably similar findings[2]. The results indicate only weak evidence for an incarceration effect, and the authors therefore conclude that the observed significant negative effect of education on teen fertility is driven primarily by the human capital effect. One potential reason for this finding could be that, in both countries, the reforms that were used for identification had a relatively minor "bite", affecting a comparatively small fraction of youth.

But other institutional features also imply variation in the amount of schooling that individuals obtain without varying the age at which they leave school. Two examples include school entry policies and changes to the length of the school day. McCrary and Royer (2011) exploit that in the US different school-entry policies imply that individuals who are

[1] A tranche of literature uses this strategy to elicit the causal impact of education on overall fertility, with diverse findings. Using variation in education induced by compulsory schooling laws in 8 European countries, Fort, Schneeweis, and Winter-Ebmer (2011) find an increase in education is associated with a large decrease in childlessness and increase in child parity, whereas Monstad, Propper, and Salvanes (2008) find no overall effect in Norway, and León (2004), using US census data, finds a decrease in the average number of children per woman. In the UK Braakmann (2011) using survey data finds a marginal increase overall fertility, whereas Clark, Geruso, and Royer (2014) using cohort-level administrative data find little effect in completed fertility.

[2] Using General Household Survey data, Silles (2011) finds effects of a similar magnitude for the UK.

born in adjacent months face different required lengths of schooling, but can effectively leave school at the same time. In terms of the mechanisms outlined in Black et al. (2008), the variation in education generated by the school entry policy should only have an impact on teenage fertility via the human capital effect. Using data from California and Texas to investigate the impact of education on a number of socio-economic outcomes they find no effect of education on the timing of fertility, which, in contrast to Black et al. (2008) suggests a relatively minor importance of the human capital effect. Berthelon and Kruger (2010) evaluate a policy intervention in Chile which increased the length of the school day. The policy had been widely criticised as previous evaluations indicated a negligible effect on educational attainment, suggesting no human capital effect. However the analysis shows that the intervention induced a significant impact on non-academic outcomes, specifically an amelioration of risk behaviours such as teen fertility and crime participation. The authors posit that this effect is entirely due to increased incarceration, as adolescents received more adult supervision per day and therefore had less time to engage in risky activities. These aforementioned studies indicate conflicting evidence with regard to the channels through which the mandatory schooling requirement influences the likelihood of early childbearing.

This paper investigates the effect of education policy on adolescent fertility in England and Wales, exploiting exogenous variation in the length of compulsory schooling induced by an institutional change, the Raising of School Leaving Age (RoSLA), implemented by the UK Government in 1972. The contribution of this paper to the literature is threefold. First, the analysis considers a legislative increase to the compulsory school leaving age which, in contrast to those studied in Black et al. (2008), impacted a significant proportion of the population. Second, as eligibility for the reform was determined by a single cut-off date the analysis proceeds using a Regression Discontinuity Design (RDD) and the paper contributes by addressing methodological concerns with implementing RDDs which have been highlighted in the recent econometrics literature. Third, as the UK has one of the highest rates of teenage pregnancy in Western Europe, the paper contributes to the body of international evidence that analyses the influence of education policy on the timing of fertility.

The analysis uses data from the Labour Force Survey, the largest representative UK household survey, exploiting an institutional change which increased the duration of compulsory education by one year. As the legislative change was implemented nationwide at a single point in time, it can be thought of as a natural experiment, which induced exogenous variation in the length of education received by an individual. The variation was determined solely by a discontinuous function of an observed covariate, the individual's month and year of birth, and therefore the estimation proceeds through a regression discontinuity design (RDD), an approach which allows the identification of causal treatment effects in quasi-experimental settings. The analysis employs both parametric and non-parametric methodologies to estimate the direct impact of the reform, and a two-stage 'fuzzy' RDD approach is used to address a pertinent policy question, namely quantifying the consequence of increasing mandatory education by one year on teenage motherhood. The results suggest that the impact of RoSLA varies non-monotonically throughout the teen years and, in contrast to Black et al. (2008), reveals strong evidence of the incarceration effect, as well as the beyond incarceration effect which may be attributable to increased human capital acquisition. The findings are robust to the empirical methodology employed and the sensitivity of the estimates to the choice of bandwidth is explored. In addition, the analysis is extended to examine the extent of the bite of the reform by investigating the extent of the impact of the treatment beyond just the teenage years, the results suggesting that RoSLA essentially caused a postponement of fertility to the late teenage years, with no observed impact of the reform after age 20.

The remainder of the paper is structured as follows: Section 2.2 summarises the institutional context. Section 2.3 describes the data used in the analysis. The econometric methodology is outlined in Section 2.4. Section 2.5 presents the results and offers interpretations, Section 2.6 concludes.

2.2 Institutional Setting

Compulsory schooling was introduced to the UK towards the end of the 19th Century, with separate rules governing school-starting and school-leaving ages. A child is required to commence education no later than the beginning of the academic year[3] after which she reaches the compulsory school-starting age of 5 years, which has remained unchanged since its inception through the Forster Education Act (1870). The first minimum school-leaving age of 10 years was introduced by the Elementary Education Act (1880), with incremental increases to the school-leaving age introduced by subsequent legislation[4].

This paper concentrates on the exogenous variation in the minimum education requirement induced by the Education (Butler) Act (1944), which initially established a minimum compulsory school-leaving age of 15. The act made provision for a further raise of the

[3] In England and Wales the academic year runs from September 1st until August 31st in the next calendar year.

[4] The Elementary (School Attendance) Act (1893) increased the age requirement to initially to 11, and up to 12 with an amendment to the act in 1899; another increase up to age 14 followed the Fisher Act (1918); the Butler Act (1944), enacted in 1947, enabled further rises first to age 15 and subsequently 16; the Education Act (2008) introduced an initial increase to age 17, and from September 2015 requires formal participation in education or training of individuals in England and Wales until their 18th birthday.

school leaving age up to age 16, but did not mandate a specific implementation date.[5] In the immediate post-war period implementation was not possible due to acute shortages in capital, material and labour, the latter so extreme that during the 1950's there were calls to reduce the length of compulsory education in order to increase the size of the labour force pool. However following the Crowther Report (1959) there was a distinct shift in attitude in favor of increasing the duration of mandatory schooling, leading to the announcement in 1964 of the government's intention to implement an increased school-leaving age in September 1970. Preparations for the age-rise were extensive and included a revised curriculum, large-scale teacher-training to increase the supply of teachers, and a building initiative enlarging schools to accommodate the increased number of students. These preparations were halted due to fiscal constraints imposed following the 1967 devaluation of sterling, with the government delaying implementation by two years. The new school-leaving age was finally introduced by Statutory Instrument 444 (1972), commonly known as the Raising of School Leaving Age (RoSLA[6]), implemented in September 1972 thus affecting academic cohorts born from 1st September 1957 onwards.

The reform impacted the leaving decisions of individuals in the lower tail of the education distribution only. Figure 2.1(a) depicts the fraction of individuals leaving education before the age of 16 by their academic cohort of birth. This proportion was steadily declining prior to the implementation of RoSLA, but there is an immediate drop of approximately 20 percentage points exactly coinciding with the introduction of the new minimum school-leaving age, indicating that the RoSLA reform constituted a binding constraint for this proportion of the school age population. Compliance with the increased mandatory age was almost ubiquitous. Since the Education Act (1962) an individual did not become eligible to leave school on the exact day he attained the compulsory school-leaving age, instead two school exit dates were imposed - the end of the Spring term (at Easter) for individuals within an academic cohort whose birthday lay between September and January, and the last day of the Summer term for those attaining school-leaving age between February and August. The implication of this 'Easter Leaving Rule' was that summer-born children born at the end of the academic year would become eligible to leave school just before the birthday where they reached compulsory school-leaving age. Specifically as the end of the Summer term usually falls around the end of June, one sixth of the first cohort directly affected by RoSLA (those born in July and August 1958), could leave school at age 15 and still be compliant with the minimum school-leaving age requirements, and therefore in Figure 2.1(a), the proportion of individuals leaving education by age 15 does not fall to exactly zero after the implementation of the increased schooling requirement.

[5]Section 35 of the Act states that the subsequent raise should occur 'as soon as the Minister is satisfied that it has become practicable to raise to sixteen the upper limit of the compulsory school age'

[6]A comprehensive history of the RoSLA can be found in Woodin, McCulloch, and Cowan (2013).

Figure 2.1: Participation in Education

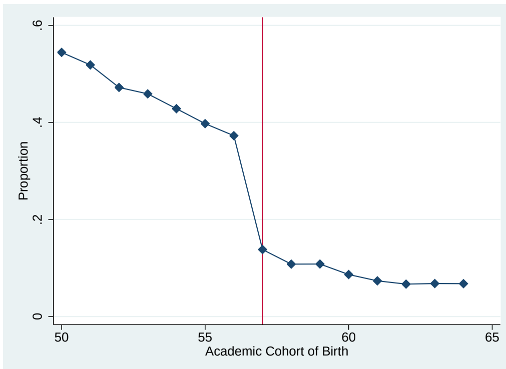

(a) Left education before age 16

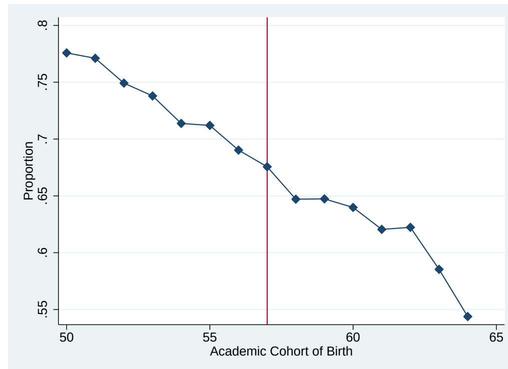

(b) Left education before age 17

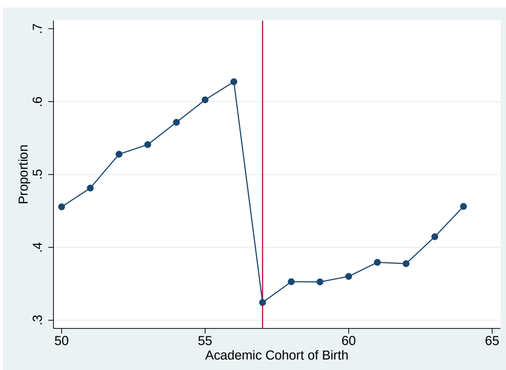

(c) Post-compulsory

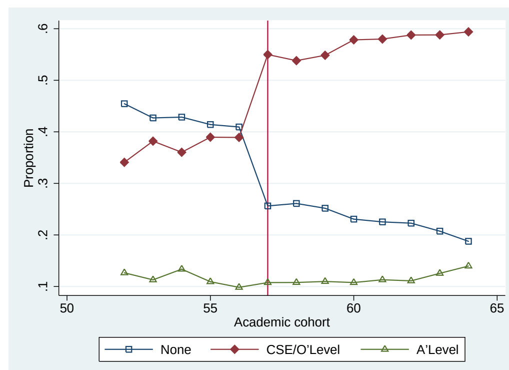

(d) Highest qualification level

Consistent with previous studies (see e.g., Chevalier, Harmon, Walker, and Zhu (2004); Dickson and Smith (2011)), the data indicate that there were no ripple-upwards effects of the RoSLA throughout the duration of education distribution. Figure 2.1(b) shows that there is no discontinuity in the downward trend of the proportion of individuals leaving education by age 17, indicating that the RoSLA did not induce an increase in the proportion of students participating in post-compulsory education. Indeed as verified in Figure 2.1(c), prior to implementation over 60% of students already participated in post-compulsory education, but approximately half of these individuals remained in school to age 16 only[7]. As a consequence it can be observed that the post-compulsory education rate actually fell approximately 30 percentage points coincidental to the introduction of the reform, afterwhich it reaches a relatively stable level consistent with the RoSLA reform inducing an increase in schooling for those individuals in the lower tail of the years

[7]The first tier of academic qualifications in England and Wales are taken at age 16, which prior to RoSLA may have been the inducement for these individuals to remain in education beyond the minimum requirement

of education distribution up to the new minimum school leaving age but not beyond. This is further supported by examining qualifications obtained: Figure 2.1(d) illustrates the trends in the highest academic qualification obtained by individuals. In the RoSLA year there is a drop in the proportion of individuals without academic qualifications of almost 15 percentage points, approximately equal to the increase in the proportion of individuals obtaining either a Certificate of Secondary Education (CSE) or Ordinary Level (O'Level) qualification, examinations which are sat in the academic year in which an individual turns 16. In contrast, there is no impact of the RoSLA on the proportion of individuals with an Advanced Level (A'Level), an examination taken at age 18.

2.3 Data

The analysis combines data from the 1975-2006 Labour Force Surveys (LFS). The survey, which is the largest representative household-level survey in the UK, contains detailed information on each individual within a household including month and year of birth, ethnicity, age at leaving full-time education, area of residence, and country of birth.

The outcome of interest in the analysis, the age at which an individual entered motherhood, is determined from the ages of the mother and the eldest child within a household at the time of the survey using the "own-children methodology" developed by Grabill and Cho (1965). This reverse-survival technique has been shown to generate age-specific fertility rates from LFS survey data which are consistent with those calculated from administrative data (Murphy and Berrington, 1993). Implicit in this procedure is that a mother-child relationship can be observed only if both individuals are present in the same household at the time of the survey. Thus in the case of parental separation the child is assumed to be resident with the mother, so that the observed mother-child relationship is biological. The determination also assumes away child mortality, and therefore the eldest child observed is primogeniture. Although these two factors may induce measurement error, it is likely that any effect would be small.[8] As the LFS contains measures of both month and year of birth it is possible to determine maternal age to within one month, a more accurate calculation than is possible with census data.[9] A further advantage of the detailed reporting of date of birth in the LFS is that it enables precise assignment of individuals to their academic

[8]The proportion of multi-family households has declined from 3% in 1961 to approx 1% in 2001 (Social Trends 32, Office of National Statistics (2002)), with over 90% of stepfamilies in 1990 being comprised of children from a previous relationship of the mother (Social Trends 38, Office of National Statistics (2008)). There has been an upward trend in single-parent families, but a fairly constant proportion of these (circa 85%) are lone-mother families (Social Trends 38, Office of National Statistics (2008)). Childhood mortality rates have been declining over time - the under-15 mortality rate stood at 31 per 100,000 in 1980, falling to 15 per 100,000 by 2000 (Child Mortality Statistics, Office of National Statistics (2010a)).

[9]For instance, the US census records year of birth only for the 1940 and 1950 censuses, thereafter also quarter of birth allowing a calculation of maternal age to within 3 months at best (Black et al., 2008).

cohort of birth, which would not be possible if only calendar year of birth was reported.

To avoid truncation of the distribution of teenage mothers, the sample is restricted to women aged between 20 and 30; the lower bound reflects that to determine whether an individual is a teen mother or not the observation must be taken after adolescence, the upper bound reflects the fact that during this period individuals started to leave the parental home from age 16 onwards, so above the age of 30 it may not be possible using information on individuals residing in a household to accurately identify whether a woman became a mother in her teenage years.

Although the LFS does report country of birth, for all but the latest surveys this measure is aggregated to the national level for UK-born individuals, as constituent countries of the UK are measured from only the 2nd quarter of 2001 onwards. Additionally, the LFS reports contemporaneous region of residence only at the time of the survey, and therefore does not have information on where an individual spent her childhood. This is problematic as the education system in Northern Ireland and Scotland differs from that in England and Wales, and in particular education in Scotland is governed by separate rules and legislation. For this reason the sample is restricted to those women who were born in the UK, but were resident in England and Wales at the time of the survey, with the implicit assumption that these individuals would have been subject to the English education system. It is therefore possible that the sample is affected by random mobility, however internal migration between constituent countries of the UK is assumed small.[10]

Table 2.1: Descriptive Statistics

Variable	Mean	Std Dev	Variable	Mean	Std Dev
Academic Cohort	57.75	4.876	Age at survey	25.38	3.063
Age left F/T Education	16.61	1.824	Subject to RoSLA	0.605	0.489
White	0.974	0.160	No of children	1.787	0.811
Mother at 15	0.003	0.053	Mother by 15	0.003	0.051
Mother at 16	0.012	0.107	Mother by 16	0.005	0.073
Mother at 17	0.028	0.166	Mother by 17	0.017	0.129
Mother at 18	0.042	0.200	Mother by 18	0.045	0.208
Mother at 19	0.050	0.218	Mother by 19	0.087	0.282
Mother at 20	0.051	0.221	Teen Mother	0.137	0.344
Number of Observations		137,502			

Table 2.1 displays the descriptive statistics for the main sample used in the analysis. The individuals were all subject to the Butler Act (1944), thus facing a minimum school-leaving age of either 15 or 16. Academic cohorts range from 1947/48 to 1964/65, with 61%

[10]Internal migration statistics are not available prior to 1991, however Stillwell, Boden, and Rees (1990) using doctor registration data from 1975-1986 estimate that the bulk of internal migration over this period was within rather than between countries of the UK.

of individuals within the sample subject to the post-RoSLA schooling regime (minimum school-leaving age of 16). The sample is predominantly white[11]; 13.7% of the sample are teenage mothers, 8.7% are mothers before the age of 19, 4.5% before age 18, 1.7% before age 17, 0.5% before age 16 and 0.3% before age 15, proportions reflective of those recorded in administrative data. Amongst mothers in the sample, the number of children per mother is 1.78. This is lower than official estimates of total fertility rates, but reflects that the sample measures fertility only up to a maximum age of 30 rather than completed fertility per woman.[12]

2.4 Empirical Methodology

As the RoSLA reform was implemented nationwide at a single point in time, it can be thought of as a natural experiment inducing exogenous variation in the length of education received by an individual. As this variation was determined solely by a discontinuous function of an observed covariate, the individual's birth date, the estimation proceeds through a regression discontinuity design (RDD), an approach which allows the identification of causal treatment effects in quasi-experimental settings. The method dates back to Thistlethwaite and Campbell (1960), who introduced the approach analyzing the impact of winning a scholarship on subsequent academic outcomes. More recently RDDs have gained popularity in applied economics and have been used to investigate, *inter alia*, the impact of impact of class sizes on scholastic achievement (Angrist and Lavy, 1999), voting shares (Lee, 2001) and labour market discrimination (Hahn, Todd, and van der Klaauw, 1999).

The RDD approach is based on the idea that a discontinuity in the assignment function to a treatment is induced in situations where individuals are deterministically assigned to the treatment based on whether the value of an observed covariate, the running variable, Z_i, falls on either side of a specific threshold value $Z_i = z^*$. The intuition is that individuals in the neighbourhood of the threshold value are identical in all other characteristics, apart from whether or not they are assigned to the treatment. Therefore by comparing individuals 'close' to the discontinuity from either side of the threshold, a causal effect of the treatment can be identified. As there is local randomization additional covariates are not necessary, but may improve precision of the estimates (Lee and Lemieux, 2010).

[11] The under-reporting of ethnic minority groups is well-known in the LFS. In an attempt to address this issue, 'boost' samples, which over-sample in areas with a high population density of under-represented groups, have been taken since 1984.

[12] The average number of children per woman in the sample is 1.18, which is comparable to cohort fertility rates in administrative data. ONS estimates of children per woman range from 0.99-1.62 for women aged 30 between birth years 1947-1983 (ONS, Cohort Fertility, England & Wales, Office of National Statistics (2010b)).

The assumptions to achieve identification in this context are hence twofold: a) that individuals are randomly selected into the RoSLA 'treatment'; b) that the timing of the introduction of RoSLA is not related to unobserved characteristics that determine teenage motherhood. Whether an individual was subject to the increased school-leaving age can be considered to be as good as randomly assigned for two reasons. Firstly, individuals are assigned to academic cohorts according to their date of birth, which cannot be perfectly controlled. Second, there is no possibility of announcement effects, whereby forward-looking parents could time the birth of their children according to RoSLA eligibility, as detailed in Section 2.2 plans to raise the school leaving age were not made public before 1964, by which time the first individuals who would be impacted by RoSLA had already been born.

Formally the RDD estimate α^{RDD} is calculated by taking the difference in the expected values of the outcome variable either side of the threshold of the observed running variable:

$$\begin{aligned}
E[\alpha^{RDD}|z] &= E[Y_1 - Y_0|Z = z^*] \\
&= \lim_{z^* \leftarrow z^+} E(y_i^1|z^*) - \lim_{z^- \rightarrow z^*} E(y_i^0|z^*) \\
&= \lim_{e \rightarrow 0} E(y_i^1|z^* + e) - \lim_{e \rightarrow 0} E(y_i^0|z^* - e) \qquad (2.1)
\end{aligned}$$

where Y_1 and Y_0 are respectively the 'treated' and 'untreated' population means; y_i^1 and y_i^0 are observations of individuals respectively to the right or the left of the discontinuity; the threshold level of the running variable is denoted $Z = z^*$. When the support of the running variable is continuous, e can be infinitely small close to the discontinuity so that the limits in (2.1) exist, and it is appropriate to use non-parametric methods in the estimation (Hahn, Todd, and van der Klaauw (2001)). As eligibility for the reform is deterministic, this representation is a 'sharp' RDD.

2.4.1 Non-parametric Estimation

The analysis uses kernel-weighted local polynomial smoothing to estimate the expectations either side of the threshold value of Z_i, with the treatment effect calculated as the difference between the predicted values calculated at the discontinuity. Although triangular kernels, by assigning larger weights to observations at the threshold in principle have better boundary properties (Fan and Gijbels, 1996), in practice kernel choice does not exert a significant impact on the magnitude of the estimates and rectangular kernels have become the *de facto* standard (Imbens and Lemieux, 2008). The order of polynomial smoothing is guided by the Bayesian Information Criterion[13]. Bootstrapped coefficients

[13] The Bayesian Information Criterion (BIC) indicated that a linear polynomial was appropriate over all outcome variables. The BIC applies a larger penalty for higher order terms than the Akaike Information

and standard errors are calculated.

The running variable in the analysis is the distance in time between an individual's birth and the implementation of the RoSLA reform. Time is clearly continuous, however a practical issue arises because the data contains only discrete measures, so that the lowest granularity that this distance can be calculated is in months. Lee and Lemieux (2010) argue that as long as the running variable, Z_i, is finely distributed the econometric complication is limited, as in practice data will always contain discrete measures (Imbens and Lemieux, 2008). In essence the concern with a discretely measured running variable is that it is not possible to allow e to become infinitely small in the neighbourhood of the discontinuity. Thus there is an irreducible gap between observations on either side of the threshold, and the casual effect of the programme is only identified with a parametric assumption regarding the assignment function (Lee and Card, 2008).

2.4.2 Parametric Estimation

Recall from equation (2.1), the estimate of interest is $E[Y_1 - Y_0|Z = z^*]$. The issue at hand with discretely measured data is that it is possible to observe $E[Y_1|Z \geq z^*]$, the outcome of the set of individuals at precisely the threshold or above who are subject to the treatment, and $E[Y_0|Z = z^* - e]$, the outcomes of the set of individuals strictly below the threshold who are not treated. With discrete Z_i, e takes on a finite number of values over the range $Z = z_j, j = (1, .., J)$, which implies that the limits in equation (2.1) do not exist. Specifically, the closest realisation below the threshold, where $z^* = z_k$, is $E[Y_0|Z = z_{k-1}]$ and therefore to predict $E[Y_0|Z = z_{k-0}]$ a parametric approach is required. As the outcome variable is binary, probit regressions are estimated using a treatment dummy, T, indicating whether the individual was subject to the RoSLA reform, and include a polynomial function of the running variable, z_j. Including interaction terms between the treatment dummy and the polynomial allows the polynomial coefficients to differ either side of the discontinuity[14].

Criterion, which proved to be less definitive, but indicated either a linear or quadratic polynomial according the outcome variable in question. As the role of the polynomial is to reflect the underlying data generating process that governs fertility, rather than fertility at a specific age as measured by the relevant outcome variable, the same order of polynomial was applied across all specifications.

[14] Although the polynomial is allowed to have different coefficients either side of the discontinuity, the same order of polynomial is applied, reflecting that the polynomial is capturing the underlying data-generating process. Lee and Lemieux (2010) note that constraining the coefficients of the polynomial to be the same on both sides of the discontinuity is inconsistent with the intuition behind the RDD approach as data from above the threshold would be used to estimate $E[Y_0|Z = z^*]$ and data from below the cutoff would be used in the calculation of $E[Y_1|Z = z^*]$. However this approach is often seen in the literature, see for example Silles (2011), as imposing this constraint will lead to more efficient estimates.

The estimation equation thus becomes:

$$Y_{ij} = \alpha_0 + \beta_0 T_{ij} + \gamma_0 P_j^l + \delta_0 (T_i \times P_j^l) + a_j + \epsilon_{ij} \tag{2.2}$$

where Y_{ij} is the outcome for individual i born at a distance of j, in months, from the relevant threshold; T_{ij} is a dummy variable indicating whether an individual born in month j was subject to the RoSLA reform, thus β_0 captures the impact of the treatment, and is hence the parametric estimate of α^{RDD}; P_j^l is a vector of polynomial functions of z_j, with $(l \in \mathbb{N})$ denoting the order of the polynomial; a_j is a specification error term that describes the difference between the true value at each z_j and the estimated polynomial function; ϵ_{ij} is an idiosyncratic error term.

The magnitude of the coefficient estimates of interest can be sensitive to the choice of polynomial in the running variable. A certain degree of smoothing may be desirable to minimise the influence of outliers and seasonality, although at a cost of deterioration in the model's fit. Higher degree polynomials follow the data more accurately, but may overstate outliers. With small bandwidths the number of higher degree polynomials is limited as J constrains the total parameters that can be estimated. The optimal order of the polynomial is again guided chosen according to the Bayesian Information Criterion. With this approach it is necessary to include more conservative standard errors to reflect modeling uncertainty. Lee and Card (2008) advocate inflating standard errors in relation to their goodness of fit statistic G,[15] and therefore (2.2) includes the specification error term a_j, which is assumed identical either side of the discontinuity and to be random and orthogonal to Z. The estimation computes robust standard errors with random, identical specification errors by clustering on z_j.

In practice both the parametric and non-parametric approaches should yield similar estimates of the RDD parameter as long as the discretisation of Z is not too coarse. Therefore Section 2.5 presents results utilising both methodologies in order to illustrate that the analysis does not rely on one particular method or specification.

[15] The Lee and Card (2008) G-statistic is calculated as:

$$G \equiv \frac{(RSS_R - RSS_{UR})/(J-K)}{RSS_{UR}/(N-J)}$$

where RSS_R is the residual sum of squares for the model using polynomial functions and RSS_{UR} for the unrestricted model using dummies respectively. Under the assumption of normality, G follows an $F_{(J-K,N-J)}$ distribution, with K the number of parameters estimated in the restricted model, N the number of observations and J the total number of values in the support of Z. The null hypothesis is that there is no systematic difference in the residual sum of squares in the restricted and unrestricted estimations.

2.4.3 Bandwidth Choice

A key issue in both the parametric and non-parametric approaches is the determination of the appropriate size of the window around the discontinuity to use in the estimation. From a theoretical perspective, by taking the limits either side of the threshold the smallest window width around the discontinuity yields unbiased estimates of the true treatment effect. However such an estimation would use only a paucity of data points and therefore have little statistical power. Wide bandwidths use a greater number of observations and will produce more efficient estimates, however a degree of bias may be introduced by including observations far from the discontinuity, the concern being that there may be unobserved changes over the bandwidth period, for instance to legislation or benefit entitlement[16], which could independently impact the proclivity toward teen motherhood, potentially confounding the analysis. It might also be expected that the magnitude of the treatment effect is different for those cohorts closer to the timing of the implementation. In addition, too great a window size may indicate a sizable treatment effect even when the data is smoothly distributed around the discontinuity. There is therefore an inherent trade-off between bias and efficiency in choosing the appropriate window of observations to include in the estimation.

Ludwig and Miller (2007) propose an optimal bandwidth selection procedure specific to a RDD context. For each candidate bandwidth, h, the cross-validation function is computed via a leave-one-out procedure, whereby for each observation, i, a regression is estimated omitting observation i and the difference is calculated between the predicted value for observation i from this regression, $\hat{y}(z_i)$, and the actual value y_i. To reflect that RDD estimates are estimated at the boundary, if the value of the running variable for observation i is to the left of the threshold, then the regression uses only observations where $z_i - h \leq z < z_i$. If observation i has a value of Z to the right of the threshold then the regression uses only observations where $z_i < z \leq z_i + h$. Repeating this procedure for each observation i with every possible bandwidth h yields the cross-validation function $CV_Y(h) = \frac{1}{N}\sum_{i=1} N(y_i - \hat{y}(z_i))^2$. The optimal bandwidth is then the value of h that minimises $CV_Y(h)$, the mean square difference of the predicted value to the true value of Y (Imbens and Lemieux, 2008).

2.4.4 Fuzzy RD

The methodology presented thus far allows the estimation of the impact of an increase in mandatory education from age 15 to age 16 on adolescent motherhood. However a more

[16] The Child Benefit Act (1975), enacted 1977, replaced family and child tax allowances paid to the household with child benefit paid directly to the primary child caretaker (usually the mother). Therefore estimates using a window width larger than 5 years may reflect the introduction of this benefit entitlement.

general determination of the impact of schooling duration on fertility behaviour may be pertinent to policy formation. As the education and fertility decisions are interrelated[17], a simple estimation of the impact of schooling on fertility using Ordinary Least Squares (OLS) may produce biased estimates.

Using an instrumental variable (IV) approach is a standard method to address such endogeneity. In the context of regression discontinuity design, the IV approach is a 'fuzzy' (FRD) regression discontinuity (Trochim, 1984). The FRD differs from the sharp design, described by (2.1), insofar that treatment assignment is not required to be a deterministic function of Z_i. Instead the *probability* of receiving treatment as a function of the running variable, $\Pr(T_i = 1|z_i)$, is discontinuous at the threshold, $Z_i = z^*$, as there are factors unobserved by the econometrician that can influence assignment to treatment, such that treatment participation is not perfectly predicted by the cohort rule. Hahn et al. (2001) argue that the FRD allows the determination of a Wald estimator even when the standard IV assumption is violated. As the estimates are applicable only to the sub-population of individuals, for whom the RoSLA reform actually induced an increase in the schooling (the 'compliers'), the estimated coefficients therefore describe a Local Average Treatment Effect (Angrist and Imbens, 1994)[18].

As in (2.2) the estimation allows for random, identical specification errors in the estimation and receive robust standard errors by clustering on z_j. The two step approach can be written as:

$$AGELEFT_{ij} = \alpha_1 + \beta_1 T_{ij} + \gamma_1 P^l_j + \delta_1 (T_i \times P^l_j) + a_{1j} + \nu_{1ij} \tag{2.3}$$

$$Y_{ij} = \alpha_2 + \xi \widehat{AGELEFT}_{ij} + \gamma_2 P^l_j + \delta_2 (T_i \times P^l_j) + a_{2j} + \nu_{2ij} \tag{2.4}$$

In the first stage (2.3), the impact of the RoSLA treatment on school-leaving age for individual i born at a distance of j months from RoSLA implementation is estimated, and then included in the second stage equation (2.4). Thus the Wald estimate, ξ^{FRD}, describes the causal effect of one year of schooling on the fertility outcome of interest Y_{ij}, and is thus equivalent to the ratio of the sharp RDD estimate from equation (2.2) and the

[17] Specifically there may be non-observed characteristics that affect both the fertility and education decision. The specification may also suffer from reverse causality: an individual with low academic attainment may choose to become a mother early. This was described by Harris, Duncan, and Boisjoly (2002) as the 'Nothing-to-lose' hypothesis, as such an individual would be likely to have poor economic opportunities regardless of the timing of her fertility. However it is also plausible that an individual who experiences early fertility may elect to curtail her education prematurely in response to motherhood.

[18] Individuals (the 'always takers') who would always stay in school until age 16 would not have been affected by the increase in school leaving age. As the reform mandated compulsory attendance, the population of 'never-takers' should not exist. Key to identification is the monotonicity assumption that RoSLA had a non-negative effect on an individual's duration of schooling, so that individuals who in absence of the reform would have remained at school after age 16 reduce their duration of education in response to the RoSLA legislation (the 'defiers') are ruled out.

first stage estimate, β_1, so that $\xi^{FRD} = \frac{\alpha^{SRD}}{\beta_1}$. This has an intuitive interpretation: as not everybody responds to the treatment, the reduced form estimate has to be multiplied by the inverse of the proportion of the affected population.

2.5 Results

In Section 2.5.1 the main results explore the impact of the RoSLA reform first over each of the individual teenage years, and also the cumulative effect over the years of adolescence. To examine the extent to which RoSLA bites, the analysis is extended by investigating the extent of any impact of the treatment beyond just the teenage years. The robustness and sensitivity of the analysis is explored in Section 2.5.2. In Section 2.5.3 the analysis is extended to examine at the policy relevant question, the impact of years of education on the timing of entry to motherhood.

2.5.1 Main Results

To illustrate the transparency of the sharp RDD approach, the results are first presented graphically. Figure 2.2(a) depicts the impact of the reform on the probability of becoming a mother at age 16, whereas the cumulative of becoming a mother before the age of 17 is shown in figure 2.2(b). These graphs are estimated using the local polynomial smoothing approach, as described in Section 2.4.1, with a bandwidth of 48 months and a smoothing polynomial of degree 1. In each case the timing of the implementation of the RoSLA reform has been normalised to 0. Appendix A displays the full set of results over each of the outcome variables.

Considering fertility at each of the individual teen years, the graphs in Appendix A are indicitive of a clear difference in fertility before and after the reform, for all but mother at age 17. As RoSLA raised the age of compulsory schooling from age 15 to age 16, the observed effect at age 16 reflects the immediate 'bite' of RoSLA and can be interpreted as the direct incarceration effect associated with the requirement to complete one year of additional schooling. At ages beyond 16, the RoSLA constraint is not binding, and therefore any observed effect cannot be attributed to incarceration alone. The graphs illustrate a non-monotonic impact of the reform over the teenage years, with negative effects for motherhood at age 16 and at age 18, a negligible effect at age 17, and positive effects at age 15 and age 19.

Analytical results are presented in Tables 2.2 and 2.3. Panel A displays results estimated using the parametric procedure as detailed in section 2.4.2, for the probability of becoming a mother at a specific year of age, or before a certain age respectively (thus teenage motherhood is defined as entering motherhood before the age of 20). The first

Figure 2.2: Graphical Results - Sharp RDD

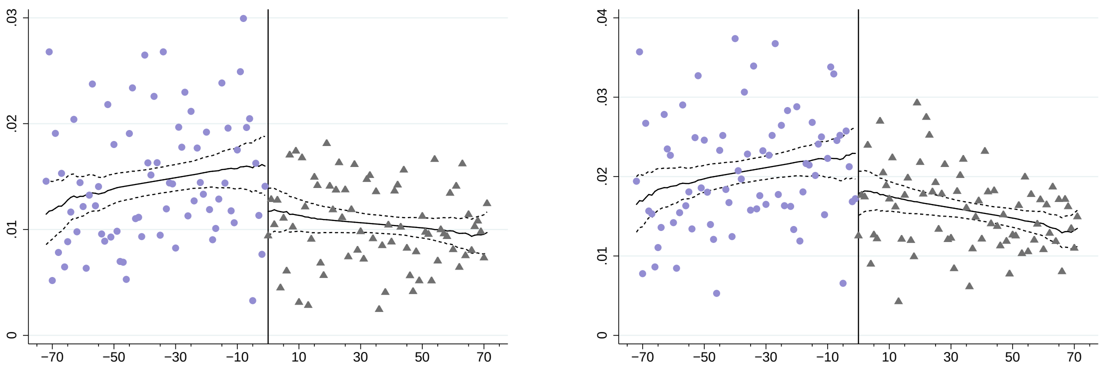

(a) Mother at age 16 (b) Mother by age 17

Notes: The graphs display local-linear polynomial smooths, as described in Section 2.4.1, using a bandwidth of 48 months, a smoothing polynomial of degree 1, and a rectangular kernel, of the probability of becoming a mother a) at age 16 and b) before age 17. The horizontal axis measures the distance, in months, of individuals' births to the RoSLA cutoff. The scatterplot indicates the proportions of mothers in each month-bin. The dashed lines are 95% confidence intervals of the local polynomial.

estimation uses the preferred bandwidth of 48 months, then estimates using half and double the preferred bandwidth are displayed to illustrate the robustness of the results to the choice of bandwidth (Imbens and Lemieux, 2008). Panel B shows the bootstrapped estimates and coefficients from the non-parametric method described in Section 2.4.1.

Examining first the estimations with fertility at a specific age as the outcome variable, the regression coefficients reveal evidence of both an 'incarceration' and a 'beyond incarceration' effect. The negative significant effect of 0.40 percentage points at age 16 reflects the direct impact of the increase in the schooling requirement, and can therefore be interpreted as the incarceration effect of RoSLA. This implies that the effect of requiring young women to stay an additional year at school is to reduce the incidence of pregnancy at the age of 16 by 36.73% relative to the sample mean. Although a positive effect at age 15 of 0.15 percentage points is observed, translating to a large increase in the incidence of pregnancy at this age, the estimate is imprecise due to the very small fraction of individuals who experience such early motherhood. A back-of-an-envelope calculation indicates approximately one quarter of the decrease in incidence of motherhood at age 16 may be attributed to individuals bringing fertility forward to age 15[19].

The pertinent question is whether the remainder of the decrease in incidence of mother-

[19]This quantitatively small effect may be attributed to individuals with preferences for extreme early fertility, who in absence of the RoSLA would have postponed motherhood to age 16, due solely to the social norm of not having a child whilst still in education. However with the increased schooling requirement these individuals find that the perceived cost of delaying fertility one more year is so great that the reform actually induces them to enter motherhood earlier than they would have done in absence of the increase in mandatory schooling.

Table 2.2: Sharp RDD - Mother at specific ages

	At 15	At 16	At 17	At 18	At 19
Panel A:					
BW = 48	0.0015	-0.0040*	-0.0019	-0.0081**	0.0058*
N = 64,359	(0.0011)	(0.0022)	(0.0026)	(0.0034)	(0.0031)
% change	49.67%	-36.73%	-6.85%	-19.63%	11.41%
BW = 24	0.0000	-0.0059*	-0.0021	-0.0038	0.0049
N = 31,566	(0.0015)	(0.0030)	(0.0035)	(0.0051)	(0.0036)
BW = 96	0.0008	-0.0047***	-0.0032	-0.0072***	0.0025
N = 124,458	(0.0008)	(0.0014)	(0.0020)	(0.0026)	(0.0025)
Panel B:					
BW = 48	0.0014*	-0.0043**	-0.0020	-0.0076**	0.0055
N=64,359	(0.0007)	(0.0018)	(0.0028)	(0.0031)	(0.0035)

Notes: Panel A displays estimates from the parametric estimations, as described in Section 2.4.2, of each dependent variable over columns, with different bandwidths over rows. Robust standard errors, which allow for random and identical specification errors, are reported in parentheses. Panel B shows bootstrapped coefficients and associated standard errors from the local-linear polynomial smoothing procedure described in Section 2.4.1, using 1,000 replications. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

hood at age 16 is due to individuals delaying fertility by one year only (a pure incarceration effect) or by more than one year. If pure incarceration only is present, then fertility should shift by one year, which would induce a positive impact of 10% at age 17. However, the coefficient in Table 2.2 suggests that there is no significant impact of the reform at age 17, in turn implying that some individuals who were not directly constrained by the RoSLA reform also delayed their fertility. This consequently should induce a positive impact at age 18, but the coefficient reveals that there is also a significant decrease in fertilty at age 18 of 0.81 percentage points, almost double the level impact seen at age 16, but implying a lesser decrease of the incidence of motherhood of 19.63% due to the larger number of individuals entering motherhood at this age. Therefore the results provide strong evidence of both incarceration and an additional downward impact of the reform on fertility that cannot be explained purely by incarceration. Furthermore, at age 19 there is a significant positive impact on fertility of 11.41%, which suggests that overall RoSLA induced a postponement of fertility to late teen years.

The estimates in Table 2.3 reflect the cumulative effect of the individual year impacts displayed in Table 2.2. The coefficient for mother by age 16 captures the impact of the RoSLA treatment on the probably of entering motherhood for all ages up to but not including the individuals 16th birthday. Thus the coefficient for mother by age 17 cumulates the ‘by 16’ effect with the ‘at 16’ effect from Table 2.2. Here the clear evidence of the incarceration effect is indicated by the coefficient on mothers by age 17, whereas the beyond incarceration effect is evident from the increasing magnitude of the coefficients

Table 2.3: Sharp RDD - Cumulative effect over teen years

	By 16	By 17	By 18	By 19	By 20
Panel A:					
BW = 48	-0.0007	-0.0048*	-0.0067*	-0.0145***	-0.0088*
N = 64,359	(0.0013)	(0.0024)	(0.0037)	(0.0046)	(0.0053)
% change	-14.52%	-29.69%	-15.18%	-17.06%	-6.57%
BW = 24	-0.0018	-0.0076**	-0.0096*	-0.0133*	-0.0082
N = 31,566	(0.0019)	(0.0036)	(0.0051)	(0.0074)	(0.0078)
BW = 96	-0.0002	-0.0049***	-0.0082***	-0.0154***	-0.0132***
N = 124,458	(0.0009)	(0.0016)	(0.0027)	(0.0036)	(0.0041)
Panel B:					
BW = 48	-0.0007	-0.0050**	-0.0070**	-0.0146***	-0.0090
N = 64,359	(0.0012)	(0.0021)	(0.0035)	(0.0046)	(0.0056)

Notes: See notes to Table 2.2.

for older teenage mothers.

To investigate the extent and duration to which the overall effect of RoSLA on fertility bites,[20] the analysis is extended to investigate fertility outcomes beyond the teenage years. In order to determine the effect on cumulative motherhood 'by' a particular age the sample must be restricted to individuals strictly above that age, that is to observe whether an individual became a mother at any age before her 25th birthday, we must observe her at age 25 or above. Table 2.4 presents the estimates of cumulative fertility by year up to by age 25. Results for each year of motherhood before age 25 use the sample of individuals aged 25-30; before age 24 use the sample of individuals aged 24-30 and so on.

The estimates in Table 2.4 confirm that the treatment exerted a significant impact over the teenage years only. The coefficients for each outcome in each of the sub-samples are consistent in sign and magnitude, displaying the same pattern of an increasing magnitudes for ages before 19, and a decrease in the size of the effect before age 20 (consistent with the positive impact at age 19 as shown in Table 2.2). After age 20 the impact of RoSLA on fertility is quantitatively small relative to the sample mean, and statistically indistinguishable from zero.

[20] Note that the analysis of the impact of RoSLA on fertility is restricted to the incidence and timing of fertility. To investigate quantum fertility requires knowledge of completed fertility, which is generally measured as the number of children per woman at age 45. However, as previously discussed, in order to accurate determine teenage motherhood it is necessary to restrict the sample to individuals aged between 20 and 30, and therefore it is not possible to investigate the impact of RoSLA on the number of children per woman. Administrative data indicates that there is no difference in completed fertility between pre-RoSLA and post-RoSLA cohorts beyond the long-run (downward) trend (ONS, Cohort Fertility, England & Wales, 2010), a result also found in the cohort analysis of Clark et al. (2014).

Table 2.4: Sharp RDD - Extended results - cumulative years

	By 16	By 17	By 18	By 19	By 20	By 21	By 22	By 23	By 24	By 25
25 - 30 sample	-0.0009	-0.0019	-0.0062	-0.0164***	-0.0132*	-0.0040	0.0011	-0.0017	-0.0052	-0.0110
N = 39,912	(0.0015)	(0.0029)	(0.0043)	(0.0058)	(0.0067)	(0.0070)	(0.0078)	(0.0088)	(0.0084)	(0.0094)
24 - 30 sample	-0.0015	-0.0036	-0.0066	-0.0151***	-0.0095	-0.0016	0.0037	-0.0012	-0.0059	
N = 45,621	(0.0014)	(0.0028)	(0.0041)	(0.0053)	(0.0061)	(0.0067)	(0.0080)	(0.0085)	(0.0080)	
23 - 30 sample	-0.0011	-0.0040	-0.0066*	-0.0150***	-0.0090	-0.0028	0.0034	-0.0014		
N = 51,164	(0.0014)	(0.0027)	(0.0039)	(0.0049)	(0.0057)	(0.0065)	(0.0082)	(0.0084)		
22 - 30 sample	-0.0007	-0.0046*	-0.0071*	-0.0163***	-0.0108*	-0.0030	0.0029			
N = 56,204	(0.0014)	(0.0027)	(0.0039)	(0.0049)	(0.0056)	(0.0064)	(0.0079)			
21 - 30 sample	-0.0008	-0.0048**	-0.0060*	-0.0154***	-0.0106*	-0.0029				
N = 61,023	(0.0013)	(0.0025)	(0.0036)	(0.0047)	(0.0055)	(0.0060)				
20 - 30 sample	-0.0007	-0.0048**	-0.0067*	-0.0145***	-0.0088*					
N = 64,359	(0.0013)	(0.0024)	(0.0037)	(0.0046)	(0.0053)					

Notes: The table shows estimates from local parametric estimations, as described in Section 2.4.2, of each dependent variable over columns, using different sub-samples over rows as indicated. Robust standard errors, which allow for random and identical specification errors, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

2.5.2 Sensitivity Analysis

Panel A from Tables 2.2 and 2.3 included estimates using three different bandwidths, the preferred, as well as double and half this bandwidth, to illustrate the robustness of the estimates to the choice of bandwidth. The preferred bandwidth was chosen according to the cross-validation procedure as described in Section 2.4.3, calculated and examined for each of the outcome variables in turn. This analysis did not yield a unique optimal bandwidth appropriate for all outcome variables, however a bandwidth between 36 and 60 months was consistently indicated.

Figure 2.3: Cross-Validation

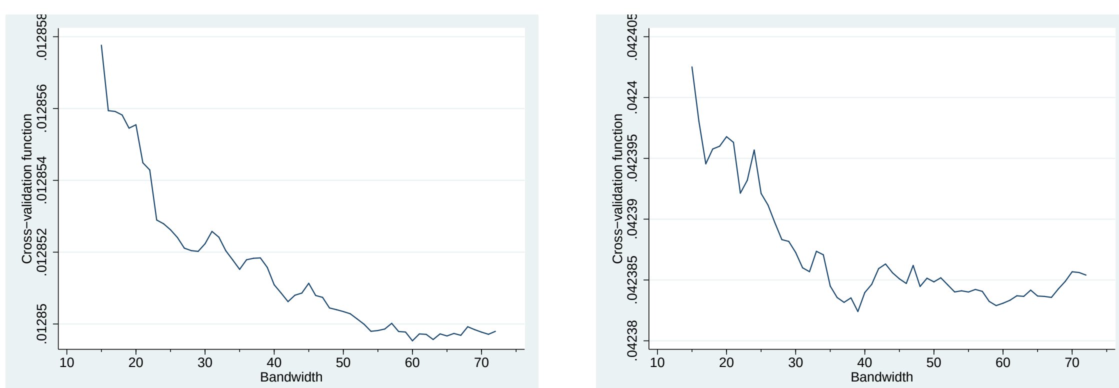

(a) Mother at age 16 (b) Mother at age 18

Notes: The graphs display the cross-validation function calculated as described in Section 2.4.3. The optimal bandwidth is given by the minimand of the function $CV_Y(h) = \frac{1}{N}\sum_{i=1} N(y_i - \hat{y}(z_i))^2$.

As an illustration, Figure 2.3(a) displays the cross-validation function for mother at age 16 over bandwidths ranging from 15 to 72 months. The function decreases in value as the size of bandwidth increases, but the graph suggests that increases in bandwidth above 40 exert little difference in the magnitude of the function. The cross-validation function for mother at age 18 is displayed in Figure 2.3(b). In this case the function does suggest a clear minimand, at approximately 40 months. The cross-validation functions for each of the outcome variables are displayed in Appendix A.

A corollary to the cross-validation procedure is to directly examine the sensitivity of the estimates to bandwidth choice. Figure 2.4 displays the magnitude of the coefficients estimated using bandwidths ranging between 18 and 72 for fertility at (a) age 16 and (b) up to, but not including age 17. The estimated impact displays some sensitivity to smaller bandwidths, but the magnitude of the estimates is essentially stable for bandwidths greater than 40. This is a reflection of what was seen in Figure 2.3(a), that increases in bandwidth exert little effect on the cross-validation function for bandwidths greater than 40. Appendix A includes the full set of results displaying the sensitivity of the estimates to bandwidth choice over each of the outcome variables. The graphs generally indicate stability in the estimated coefficients for all outcome variables at bandwidths from approximately 40 onwards, apart from the estimates for mother at age 18 (also affecting cumulative fertilty by ages 19 and 20), which achieve stability after approximately 60 months.

Figure 2.4: Sensitivity of Estimates to bandwidth choice

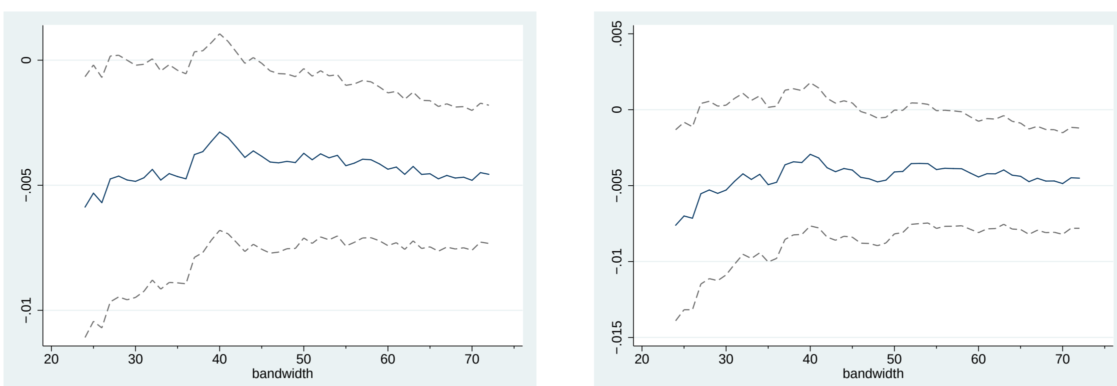

(a) Mother at age 16 (b) Mother by age 17

Notes: The graphs display the magnitude of the estimates, along with the 95% confidence interval, over different bandwidths based on the parametric regression discontinuity design as described in Section 2.4.2.

At the boundary the comparison is between individuals born at the end (August) of one academic cohort with individuals who are born at the beginning (September) of the next academic cohort. The key identifying assumption is that individuals in the neighbourhood of the discontinuity are identical in characteristics apart from their assignment to the treatment. However there may be fundamental differences in individuals according to their relative and social age within an academic cohort and therefore the RDD estimation, which is essentially a between-cohort comparison at the boundary, may just reflect compositional differences of those born at the beginning versus the end of a cohort. For instance, Crawford, Dearden, and Meghir (2010) find that relative age within a cohort exerts an important influence on academic outcomes, younger individuals in a cohort perform on average significantly worse than their older peers in assessments, which the authors attribute to the absolute age of the individual when taking the test. In the context of fertility behaviour, *a priori* it may be expected that older individuals within a cohort would have higher fertility due to their higher emotional and physical maturity, as forging a relationship requires a set of social skills that are likely to be more developed in individuals born earlier within a cohort. In addition because fecundability increases over the period of adolescence (Wood and Weinstein, 1988), older individuals are more able to conceive. However, analyzing the fertility outcomes within academic cohorts in Sweden, Skirbekk, Kohler, and Prskawetz (2004) find that individuals born at the beginning of a cohort actually enter motherhood up to 4.9 months later than those born at the end of the academic cohort, which the authors attribute to the 'social age' effect.

Table 2.5: Placebo Analysis

Panel A	At 15	At 16	At 17	At 18	At 19
1951	-0.0032**	0.0011	0.0002	0.0010	-0.0005
N=42,803	(0.0014)	(0.0026)	(0.0031)	(0.0043)	(0.0043)
RoSLA	0.0015	-0.0040*	-0.0019	-0.0081**	0.0058*
N=64,359	(0.0011)	(0.0022)	(0.0026)	(0.0034)	(0.0031)
1964	-0.0023**	0.0013	-0.0032	0.0000	-0.0032
N=73,021	(0.0009)	(0.0017)	(0.0025)	(0.0033)	(0.0034)

Panel B	By 16	By 17	By 18	By 19	By 20
1951	-0.0019	-0.0009	-0.0006	0.0004	-0.0002
N=42,803	(0.0025)	(0.0040)	(0.0035)	(0.0068)	(0.0106)
RoSLA	-0.0008	-0.0049**	-0.0068*	-0.0147***	-0.0089*
N=64,359	(0.0013)	(0.0024)	(0.0037)	(0.0046)	(0.0053)
1964	-0.0016	-0.0002	-0.0034	-0.0033	-0.0066
N=73,021	(0.0012)	(0.0020)	(0.0036)	(0.0040)	(0.0050)

Notes: The table shows estimates from parametric estimations, as described in Section 2.4.2, of each dependent variable over columns using the preferred bandwidth of 48 months, with the discontinuity defined in different years over rows. Robust standard errors, which allow for random and identical specification errors, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

In order to confirm that the results presented in Section 2.5.1 are indeed driven by the reform rather than inherent between cohort effects two further robustness checks are undertaken. Firstly a falsification exercise is undertaken, placebo regressions are estimated under the assumption that RoSLA was implemented prior to or after actual implementation. The results of the placebo analysis are displayed in Table 2.5, which are not consistent with the estimates that use the correct RoSLA assignment. The sign, magnitude and significance of the coefficients differ non-systematically, suggesting that observed effect on fertility is in fact driven by the implementation of RoSLA.

Table 2.6: RD-DiD estimates

Panel A	At 15	At 16	At 17	At 18	At 19
Pre-RoSLA DiD	0.0013	-0.0066***	-0.0116***	-0.0188***	0.0006
N = 79,852	(0.0010)	(0.0017)	(0.0029)	(0.0036)	(0.0035)
Pre-Post RD-DiD	0.0008	-0.0043***	-0.0037*	-0.0091***	0.0007
N = 137,502	(0.0008)	(0.0012)	(0.0022)	(0.0025)	(0.0028)

Panel B	By 16	By 17	By 18	By 19	By 20
Pre-RoSLA DiD	0.0003	-0.0062***	-0.0179***	-0.0363***	-0.0355***
N = 79,852	(0.0012)	(0.0022)	(0.0037)	(0.0051)	(0.0061)
Pre-Post RD-DiD	0.0000	-0.0043***	-0.0080***	-0.0169***	-0.0165***
N = 137,502	(0.0009)	(0.0015)	(0.0026)	(0.0036)	(0.0042)

Notes: The table shows estimates from the Regression Discontinuity difference in difference procedure, as described in Section 2.5.2 of each dependent variable over columns, using non-overlapping windows of observations and a bandwidth of 36 months. The Pre-RoSLA RD-DiD is estimated over the 47/48 - 52/53 and 53/54 - 59/60 windows. The Post-RoSLA RD-DiD is estimated over the 53/54 - 59/60 and 60/61 - 64/65 windows. The Pre-Post RD-DiD is estimated over all three windows. Robust standard errors, which allow for random and identical specification errors, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Second, following Danzer and Lavy (2013), a difference in difference approach is applied in the context of the regression discontinuity design (RD-DID). This procedure explicitly nets out any inherent between cohort differences at the August-September threshold by using three non-overlapping windows of observations[21]- the pre-RoSLA period (academic cohorts 1947/48 - 1952/53), the post-RoSLA period (1960/61 - 1964/64) and the period around the RoSLA discontinuity (1953/54 - 1959/60). For each sub-period the running variable is defined as the distance in months from the relevant August-September threshold. Two versions of following specification are then estimated:

$$Y_{ij} = \beta_0 + \beta_1 Right_{ij} + \beta_2 RoslaRight_{ij} + \Sigma_{k=1}^{3} Period_k + \gamma_0 P_j^l + \delta_0 (T_i \times P_j^l) + a_j + \epsilon_{ij} \quad (2.5)$$

[21] Distinct windows are required to form the counterfactual observations. In order to accommodate the total observation window a bandwidth of 36 months is used in the estimations

where Y_{ij} is the outcome of interest for individual i born at a distance of j from the relevant threshold; $Right$ is an indicator variable for an observation being on the right-hand side of the relevant discontinuity; $Period$ are period dummies for each window of observations; $RoslaRight$ is a dummy equal to 1 if the observation is on the right-hand side of the discontinuity in the period around the RoSLA discontinuity, thus β_2 describes the RD-DiD estimate. The γ and δ capture the polynomial smooth in the running variable.

Table 2.6 presents the results of the difference-in-difference analysis considering first the pre-RoSLA period only as counterfactual observations, and second using both pre and post-RoSLA periods for comparison. The estimates are qualitatively similar to those presented in the main analysis and therefore adjusting the original RoSLA coefficients to account for any inherent between-cohort discontinuities does not an induce a significant impact on the sign or magnitude of the RDD estimates.

2.5.3 Further Estimations

Figure 2.5: Graphical results - Fuzzy RDD

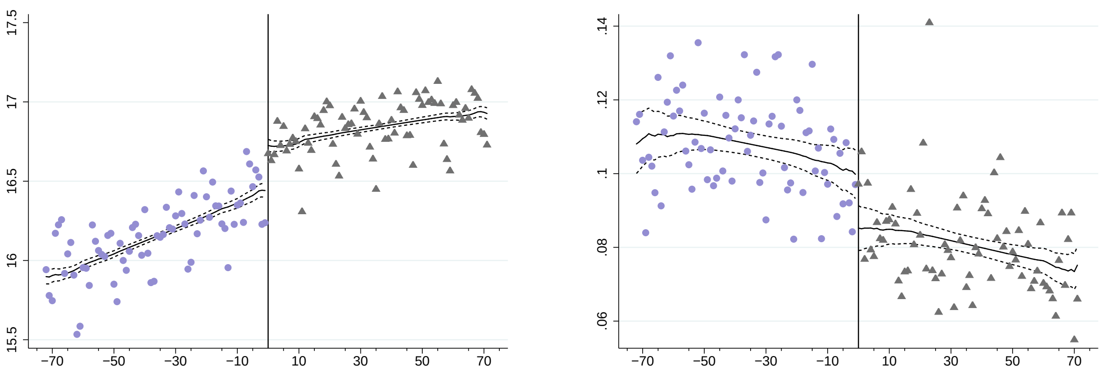

(a) Age left full-time education (b) Mother by age 19

Notes: The graphs display local-linear polynomial smooths, as described in Section 2.4.1, using a bandwidth of 48 months, a smoothing polynomial of degree 1, and a rectangular kernel, for a) age an individual left school (first-stage of the fuzzy RDD) and b) the probability of becoming a mother before age 19 (second-stage of the fuzzy RDD) . The horizontal axis measures the distance, in months, of individuals' births to the RoSLA cutoff. The scatterplot indicates the proportions of mothers in each month-bin. The dashed lines are 95% confidence intervals of the local polynomial.

Finally the analysis considers the impact of education as measured by years of schooling on adolescent fertility in a two-stage approach. In the first stage the impact of the RoSLA reform on schooling duration, measured by the age at which an individual finished full-time education is measured. This prediction is used in the second stage to analyse the effect on the probability of entry to motherhood. Figure 2.5 presents these two stages graphically.

The analytical results are reported in Tables 2.7 and 2.8. The top panel presents the Wald Estimates using the preferred bandwidth of 48 months, as well as estimates produced using half and double the preferred bandwidth. The middle panel displays results of simple OLS regressions of the impact of years of schooling on the probability of teen motherhood, and the bottom panel presents the reduced form and first stage of the estimation (for expositional convenience only the preferred bandwidth estimates are reported in these latter panels).

Table 2.7: Fuzzy RDD - Impact of years of education - individual years

	At 15	At 16	At 17	At 18	At 19
Wald Estimates					
BW = 48	0.0047	-0.0147	-0.0064	-0.0250*	0.0202*
N= 64,359	(0.0035)	(0.0082)	(0.0086)	(0.0124)	(0.0100)
BW = 24	-0.0010	-0.0207	-0.0079	-0.0126	0.0179
N = 31,566	(0.0050)	(0.0124)	(0.0120)	(0.0184)	(0.0126)
BW = 96	0.0024	-0.0176**	-0.0132*	-0.0244**	0.0091
N = 124,458	(0.0025)	(0.0054)	(0.0065)	(0.0092)	(0.0081)
OLS					
Years of Education	-0.0004**	-0.0030***	-0.0069***	-0.0088***	-0.0096***
	(0.0001)	(0.0002)	(0.0002)	(0.0003)	(0.0003)
IV					
Reduced Form	0.0014	-0.0043*	-0.0019	-0.0073**	0.0059*
	(0.0010)	(0.0023)	(0.0027)	(0.0033)	(0.0030)
First stage	0.2934***	0.2934***	0.2934***	0.2934***	0.2934***
	(0.0586)	(0.0586)	(0.0586)	(0.0586)	(0.0586)

Notes: The table shows estimates from local parametric estimations, as described in Section 2.4.2, of each dependent variable over columns, using a bandwidth of 48 months. First-stage F-statistic = 25.07. Robust standard errors, which allow for random and identical specification errors, are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The OLS coefficients consistently indicate that there is a negative relationship between an individual's propensity of early motherhood and the age at which she left full-time education. However, as discussed in Section 2.4.4, there may be omitted variables which imply that the residual term is correlated with years of education. If the unobserved heterogeneity is such that it asserts a positive impact on the propensity for early motherhood and a negative impact on schooling, then the OLS coefficients will be biased upwards. Conversely if the unobserved heterogeneity impacts both teen motherhood and years of schooling in the same direction, then the OLS estimates will be understated. This potential endogeneity is addressed using the FRD procedure described in Section 2.4.4. Recall this is analogous to an IV approach, where the RoSLA treatment is applied as an instrument for schooling. The identification assumption is that the timing of the RoSLA implementation is orthogonal to unobserved determinants of motherhood, and therefore the effect of the reform on fertility can be understood as operating only through its impact on years of education. The first stage reveals that the reform had a significant positive impact on years of schooling, raising it on average by approximately 3 months, which reflects that prior to implementation of RoSLA a substantial proportion of the school age population already stayed at school until at least age 16, as depicted in Figure 2.1(c). Considering the wald estimates over the individual years, Table 2.7, of the effect of the duration of education on teen motherhood, these differ from the OLS estimates non-systematically: the coefficients on mother at age 15 and mother at age 19 change sign (from positive to negative) indicating that the OLS estimates of these coefficients are downwardly biased. The coefficients on mother at age 16 and at age 18 are the same sign (negative) as the OLS coefficients, and are larger in magnitude indicating that the OLS estimates are understated. The coefficient on mother at age 17 also has the same sign (negative) but is smaller in magnitude than the OLS coefficient. These observations imply that not only does RoSLA have a varying impact on fertility depending on the age of the mother, but also that the correlation between unobserved factors and years of schooling varies throughout the teen years.

Table 2.8: Fuzzy RDD - Impact of years of education - cumulative years

	By 16	By 17	By 18	By 19	By 20
Wald Estimates					
BW = 48	-0.0026	-0.0173	-0.0237	-0.0487**	-0.0226
N= 64,359	(0.0047)	(0.0094)	(0.0122)	(0.0164)	(0.0187)
BW = 24	-0.0058	-0.0248	-0.0296	-0.0398	-0.0207
N = 30,338	(0.0065)	(0.0150)	(0.0188)	(0.0269)	(0.0297)
BW = 96	-0.0011	-0.0193**	-0.0346***	-0.0575***	-0.0519**
N = 118,388	(0.0031)	(0.0062)	(0.0098)	(0.0132)	(0.0161)
OLS					
Years of Education	-0.0007***	-0.0037***	-0.0106***	-0.0194***	-0.0290***
	(0.0001)	(0.0002)	(0.0003)	(0.0004)	(0.0006)
IV					
Reduced Form	-0.0008	-0.0051*	-0.0070*	-0.0143**	-0.0084
	(0.0014)	(0.0026)	(0.0038)	(0.0045)	(0.0053)
First stage	0.2934***	0.2934***	0.2934***	0.2934***	0.2934***
	(0.0586)	(0.0586)	(0.0586)	(0.0586)	(0.0586)

Notes: see notes for Table 2.7. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Considering the cumulative fertility estimates, Table 2.8, the first stage of course is identical to that in Table 2.7. A comparison of the OLS and the wald estimates reveals that they all share the same sign, in contrast to the results in Table 2.7 for the estimations at each of the individual teen ages. The cumulative estimates thus suggest that any positive correlation between the measure of education and unobservables (such as at age 16) is offset by negative correlation (for instance at age 15).

To reconcile the differences between the reduced form (sharp RDD) and the wald (fuzzy RDD) estimates, recall that the SRD measures the causal effect of the reform, which is the average effect of being subject to the RoSLA regime in comparison to the pre-RoSLA regime (on average an extra three months of schooling). In contrast, the FRD approach rescales the reduced form results so that the Wald estimates reflect the effect on the propensity for motherhood of an additional year of education for the sub-population of individuals who were induced to increase the duration of schooling by the RoSLA reform.

2.6 Conclusion

This paper has investigated the impact of an increase in the minimum compulsory school-leaving age on teenage fertility rates, using data from the UK Labour Force Survey, the largest representative UK household survey. The findings indicate a non-monotonic impact over the individual teenage years. In contrast to previous research, the results provide strong evidence of a large incarceration effect. This discrepancy may be explained by the proportion of individuals directly affected by the institutional change to mandatory education. The Norwegian reform analysed by Black et al. (2008) increased the duration of schooling by two years, yet the estimated increase to individuals' education was just 0.122 years, indicating that only a small fraction of the population were impacted. In contrast the UK's RoSLA, compelling an increase to compulsory schooling of just one year, increased the average years of schooling by 0.293 years due to the higher proportion of individuals affected. Hence although the incarceration effect, by capturing the shift in fertility for the age at which the legislation bites, may be thought of as just a mechanical response to the extra year of schooling induced by the legislation change, the evidence suggests that if mandating a higher school graduating age raises the schooling durations of a large share of the school-age population, teenage fertility rates will be substantially affected.

Unfortunately, the data used in this analysis does not allow examination of the mechanism that results in the beyond incarceration effect, the question therefore remains to what extent this is attributable to the impact of education on human capital acquisition. Extending the analysis beyond the teenage years revealed that the impact of RoSLA was to essentially induce a postponement of fertility from early teen to the late teenage years, with a large increase in the incidence of fertility at age 19, and the impact of the increase in compulsory education tailing off after age 20. Given that these individuals continued to bear children at a relatively young age, a question for future research is whether this postponement of fertility positively impacted outcomes for these mothers and their children.

Chapter 3

Post-compulsory education incentives and non-educational outcomes

The best defence against social exclusion is having a job, and the best way to get a job is to have a good education, with the right training and experience.

(Prime Minister Tony Blair, 1999)

3.1 Introduction

In 1999 the Connexions Strategy was launched by the UK Government to address concerns that an increasing proportion of the population was becoming detached from society, specifically that skill deficits and associated worklessness create a welfare burden and social disharmony. A key policy initiative became the expansion of education opportunities for young people, to facilitate the acquisition of requisite skills to ensure successful entry to the labour market. Of primary concern was the relatively low post-compulsory education participation rate of individuals from lower socio-economic groups, with the perception that precisely these individuals tend to have higher rates of societal problems such as substance abuse, crime and teenage pregnancy. This group of individuals in particular may face financial impediments, or the opportunity cost of forgoing immediate employment may be such that continuing education beyond the mandatory requirement is not a viable option (Rees, Williamson, and Istance, 1996).

This study examines the impact of the Education Maintenance Allowance (EMA), an intervention combining elements of a conditional cash transfer programme and a financial incentives scheme, implemented in the United Kingdom[1], on non-educational outcomes, specifically teenage fertility and crime rates. The EMA provided a financial inducement, available for a maximum of two years, for individuals remaining in education after the compulsory school-leaving age, which was means-tested on the basis of familial income, in order to encourage participation in further education and improve the academic achievement and vocational skill levels amongst young people from low-income families. The findings indicate that the programme led to strong increases in post-compulsory participation among eligible males in both years of programme availability, and for eligible females in the second year of post-compulsory schooling, which induced a significant decrease in the underage maternity rate. Weaker, but suggestive evidence is presented of a decrease in youth offending rates, predominately driven by the increase in participation by young men.

The paper contributes to the literature in three ways. First, the analysis considers the impact of a conditional cash transfer programme in a non-developing country setting, and is also related to the literature regarding the ability of financial incentives to induce educational improvements. Secondly, this paper adds to the literature on the capability of education policy to exert influence on non-educational outcomes. Third, this is the first paper to evaluate the EMA programme outside of the initial pilot period, and to specifically examine some key non-academic outcomes that the programme was implemented to influence.

Conditional cash transfer (CCT) programmes offer monetary incentives that are paid to eligible individuals contingent on specified behaviours. They have become a feature in many developing countries to encourage investments in human capital (Rawlings and Rubio, 2005), where outcomes such as nutrition, health and education are often targeted. A number of studies have documented the influence of CCT programmes on risky adolescent behaviours. Gutiérrez, Bautista, Gertler, Hernández, and Bertozzi (2004) find that the *Oportunidades* programme in Mexico to encourage increased engagement with healthcare services led to a reduction in smoking and alcohol use amongst adolescents, but the programme had no impact on the incidence of condom use. With regard to schooling programmes and fertility, Handa, Halpern, Pettifor, and Thirumurthy (2014) find that the Kenya Cash Transfer for Orphans and Vulnerable Children resulted in a decrease of 6 percentage points in sexual activity, and a 4 percentage point higher use of condoms. Baird, Chirwa, McIntosh, and Özler (2010) analyse a randomised control trial in Malawi. The programme provided cash transfers conditional on school enrollment, finding that self-reported sexual activity decreased by 38%, with the incidence of pregnancy declining

[1] The United Kingdom is comprised of four constituent countries: England, Northern Ireland, Scotland and Wales. Due to data availability the analysis is restricted to England, the largest constituent country, only.

by 30%. These latter studies provide evidence that CCT programmes to increase school attendance have a spillover effect on adolescent fertility.

CCT programmes are gaining popularity in more developed economies. Programmes to encourage participation in secondary education include the Macedonian CCT for Secondary School Education, New York's Opportunity NYC, and the Australian AUSTUDY programme. The former provided household cash transfers, conditional on school enrollment of children. Armand (2013) finds that households where payments are directed to the mother have a 9.8 percentage point higher probability of the children being enrolled or having completed secondary school. Evaluations of Opportunity NYC, where payments for school attendance were shared between parents and the child, find an improvement in outcomes amongst high school students, but no measured impact for younger pupils (Riccio, Dechausay, Greenberg, Miller, Rucks, and Verma, 2010). In contrast the AUSTUDY provided financial support directly to young people completing their final two years of education. Dearden and Heath (1996) find the programme induced a significant increase in participation rates.

Programmes providing financial incentives to students as a reward for attaining specific educational goals have mixed results. Leuven, Oosterbeek, and van der Klaauw (2010) find that achievement based rewards elicit positive improvements amongst high-ability students only, but discourage students of lower ability, whose long-term performance is negatively affected. In contrast, Angrist and Lavy (2009) find cash awards for low achieving students elicit positive results, which are predominantly driven by students at the 'margin' of certification. Furthermore the programme induced substantial long-run effects on the likelihood of enrollment in higher education. In school-based randomised trials conducted in a wide group of schools Fryer (2011) finds that student incentives are more successful at increasing achievement when rewards are conditional on inputs to the education production function rather than explicitly tied to specific attainment levels, a finding supported in Barrow and Rouse (2013), where students significantly increase time and effort devoted to educational activities as a result of incentives to meet a combination of benchmarks.

To summarise, the existing evidence suggests that both CCT and financial incentive programmes have the potential to encourage positive investments in education; that the group of students targeted by the programme, the programme goals and the payment structure are all important determinants a programme's degree of success. The EMA programme incorporates elements of both a financial incentive and a CCT programme, where a weekly disbursement is contingent on attendance only and the payment for achievement, available at the end of each term, is designed to reward engagement with, or completion of, a course. These two aspects potentially address discouragement of low-ability students.

The question remains, thus the focus of this paper, is whether the EMA programme by incentivising participation amongst individuals from low-income families, can affect adolescent risk behaviours.

Early childbearing is an issue which has attracted considerable attention in the literature. Teenage mothers tend to be characterised by low levels of academic attainment (An, Haveman, and Wolfe, 1993; Kiernan, 1997), and are more likely to come from economically disadvantaged or instable families than women who delay childbearing. Ermisch and Pevalin (2003) finds that the key determinant of teenage motherhood is prior familial pecuniary conditions, reflecting that adolescent motherhood may be the consequence of childhood economic disadvantage. Therefore the group of individuals targeted by the EMA programme also contains the specific set of young women with a higher propensity of early fertility.

Increasing the mandatory education requirement has been found to influence the incidence of teenage fertility. Black et al. (2008) relate exogenous changes in the compulsory school-leaving age to a lower incidence of teenage motherhood, proposing two mechanisms through which more schooling leads to to reduced fertility: firstly individuals are 'incarcerated' for a longer period in formal education, which may preclude opportunities to participate in risky behaviours. Second, the increased investment in human capital may raise the opportunity cost of early fertility. Insofar that the EMA incentivised voluntary rather than mandatory participation in education, the former mechanism is not directly applicable. However, evaluating a Swedish education policy which prolonged the length of post-compulsory schooling, Grönqvist and Hall (2013) find that although the increased educational requirement was not mandatory, the availability of a more academic programme of longer duration induced a significant decrease in early childbearing for individuals who chose this option.

The relationship between schooling and crime has been well documented. Using US data, Lochner and Moretti (2004) find that increased schooling reduces an individuals likelihood of arrest and incarceration, which they attribute to reduced criminal activity rather than a decrease in prosecution rates. Machin, Marie, and Vujić (2011) exploit exogenous variation in the length of schooling induced by a UK legislative change. Their results indicate significant decreases in property crime driven by the reduction in the proportion of individuals leaving school without an educational qualification. Regarding youth crime, the evidence of the effect of schooling on juvenile crime rates is mixed. For instance Jacob and Lefgren (2003) find that on school days there is a decreased rate of property crimes, but that the incidence of violent crimes actually increases. In contrast, Berthelon and Kruger (2010) evaluating a Chilean school reform that increased the number of hours spent in school per day, find a significant reduction in the adolescent crime rate. Chioda, De Mello, and Soares (2012) examine the impact on crime of *Bolsa Família*, a CCT programme implemented in São Paolo, Brazil. The authors exploit exogenous variation in programme participation induced through the expansion of the programme to adolescents aged 16 and 17, finding evidence of a substantial decrease in youth crime of 21%, driven by changes in peer interactions.

The remainder of the paper is structured as follows: Section 3.2 describes the Education Maintenance Allowance programme in detail. Section 3.3 outlines the data used in the analysis. The econometric methodology is summarised in section 3.4. Section 3.5 presents the results and offers interpretations, section 3.6 concludes.

3.2 Institutional Setting

3.2.1 Education System

Compulsory education has been a fundamental feature of the English education system for over a century. The academic year begins on September 1st and runs to August 31st. A child is required to commence school no later than the beginning of the academic year after which she turns 5, and must remain in full-time education until the legislated school-leaving age[2]. The first tier of academic qualifications, the General Certificate of Secondary Education examinations (GCSEs), are completed by June 30th in the final compulsory academic year. Conditional on performance in these examinations, those individuals opting to remain in education choose between academic courses leading to the second tier of academic qualifications, Advanced Levels (A'Levels) or enrolling in vocational courses.

After completing compulsory education, young people can choose to continue with education on a voluntary basis at Further Education Institutions (specialist colleges or schools), work-based learning (government or private training schemes tied to employment), or employment. Outside these options individuals are classified as 'Not in Education, Employment or Training' (NEET). During the late 1990's approximately 10% of the 16-18 population was classified as NEET, and the trend was on an upward trajectory. Furthermore a fundamental concern arose regarding the skill levels of school-leavers, as the proportion of 18 year olds remaining in education was 20% lower than the EU average, and adult illiteracy rates were amongst the highest in Europe (Social Exclusion Unit, 2001).

[2]During the period of analysis considered in this paper individuals in England reached the legislated school-leaving age on June 30th in the academic year they turn 16. Following the Education Act (2008) the school-leaving age has been raised up to age 17 (18) for those individuals eligible to leave school after Summer 2013 (2015), which is outside the period of analysis considered by this study.

3.2.2 The Education Maintenance Allowance

The Education Maintenance Allowance (EMA) was a conditional cash transfer programme implemented by the UK Government with the aim of increasing the participation rate in post-compulsory education by young people from low-income households. In England the programme ran for a total of 12 years, Figure B.1 in Appendix B presents a timeline of key features of the programme.

Prior to a national rollout, the intervention was subject to an extensive pilot period. The first pilots were launched in September 1999 in 15 Local Authorities (LAs) in England and one council area in Scotland. The areas chosen for the pilot were not random, rather areas with lower than average post-compulsory participation and educational attainment were chosen for the initial phase. Control areas, similar in attributes to the pilot areas, were selected to facilitate an evaluation. The pilot was extended to a further 41 English LAs in September 2000, and to three more Scottish areas in 2002[3]. Reports of the quantitative evaluation of the first pilot areas in the first and second year of the programme were released in March 2001 and 2002 respectively, which indicated significant positive impacts of the intervention on participation of individuals eligible for receipt. In response to these initial findings, the UK Government in July 2002 announced its intention to rollout the programme nationally from September 2004.

Eligibility for the allowance was determined by the familial (household) income of the individual. In the initial phase individuals living in households where total income, net of government transfers, was less than £13,000 were eligible for the maximum EMA payment of £30 a week. The minimum payment was set at £5 per week for individuals in households where net income was not higher than £30,000, and a linear taper was applied to set the level of EMA payment for those in households between the minimum and maximum thresholds. Relative to the median household income level in the first year of the pilot, £15,400, this implied that a substantial proportion of the population were eligible for the allowance. In addition to the weekly payment, all individuals in receipt of EMA were eligible to receive termly bonuses of up to £100, depending on their academic achievement and attendance record. The benefit could normally be claimed for

[3]The areas in the first phase of the pilot were Bolton, Nottingham, Cornwall, Doncaster, East Ayrshire, Gateshead, Greenwich, Lambeth, Leeds, Lewisham, Middlesbrough, Oldham, Southampton, Southwark, Stoke-on-Trent and Walsall. Control areas, for the evaluation of the initial phase, were Blackburn-with-Darwen, Blackpool, Derby, Devon, Newcastle-upon-Tyne, Norfolk, Portsmouth, Redcar and Cleveland, Rochdale, Rotherham, Stockton-on-Tees. In the second phase the pilot was extended to Barking and Dagenham, Barnsley, Birmingham, Bradford, Brent, Camden, Kingston-upon-Hull, Coventry, Ealing, Hackney, Halton, Hammersmith and Fulham, Haringey, Hartlepool, Islington, Knowsley, Lancashire, Leicester, Liverpool, Luton, Manchester, Newham, North East Lincolnshire, North Tyneside, Northumberland, Salford, Sandwell, Sheffield, South Tyneside, St Helens, Suffolk, Sunderland, Tameside, Tower Hamlets, Wakefield, Waltham Forest, Wandsworth, Wigan, Wirral, Wolverhampton, Worcestershire. The final pilot areas were Dundee, Glasgow and West Dunbartonshire.

up to two years, and could be used to attend either academic or vocational full-time post-16 education. These income thresholds and payment amounts were kept fixed in nominal terms, implying that the generosity of the programme decreased over time, as not only the real value of the weekly payments and term bonuses, but also the proportion of individuals eligible for receipt decreased with the duration of the pilot programme.

With the national rollout the threshold level for the maximum payment was increased to £19,630, the minimum payment was increased to £10 a week and the taper was replaced with an intermediate payment of £20 a week, for those individuals whose household income fell between the two thresholds. These changes imply that the programme under the national rollout was more generous, both in terms of the payments and the proportion of individuals eligible for the maximum award, but less generous in terms of the value of the maximum weekly disbursement, as this was held fixed at the nominal value of £30. Further increases to the eligibility thresholds occurred over the two subsequent years[4]. As in the pilot period, after the national rollout the income thresholds and payment amounts were not indexed to inflation, implying that overall the generosity of the programme decreased over time in both dimensions, as not only the real value of the allowance, but also that the proportion of individuals eligible for receipt decreased with the duration of the programme[5].

In October 2010, as part of the austerity measures introduced in the wake of the Great Recession, the UK Government announced its intention to withdraw the EMA from England. The programme was closed to new applicants in January 2011, and the final payments were disbursed at the end of the academic year 2011/12. The devolved Governments of Scotland and Wales have retained a scaled down version of the EMA, with a single income threshold at the maximum EMA payment level, and have eliminated the achievement and retention bonuses.

Figure 3.1 displays the proportion of the area level population, aged 16-17 at the start of the academic year, who were in receipt of the EMA by local authority pilot status for England. Given that the overall generosity of the programme decreased over time, it could be expected that the highest take up should occur within the first years of the programme, and decline thereafter. However, during the first pilot year take-up of the allowance was relatively low, reflecting that the programme was not announced until April 1999 by which time the majority of students in their final compulsory year would have already made choices regarding post-compulsory participation in September.

[4] The threshold for the maximum payment was increased in September 2005 to £20,270. Both thresholds were increased for a final time in September 2006 to £20,817 and £30,810 for the maximum and minimum payments respectively.

[5] Table B.1 in the Appendix compares the income thresholds to median income for the duration of the programme.

Figure 3.1: EMA Take-up

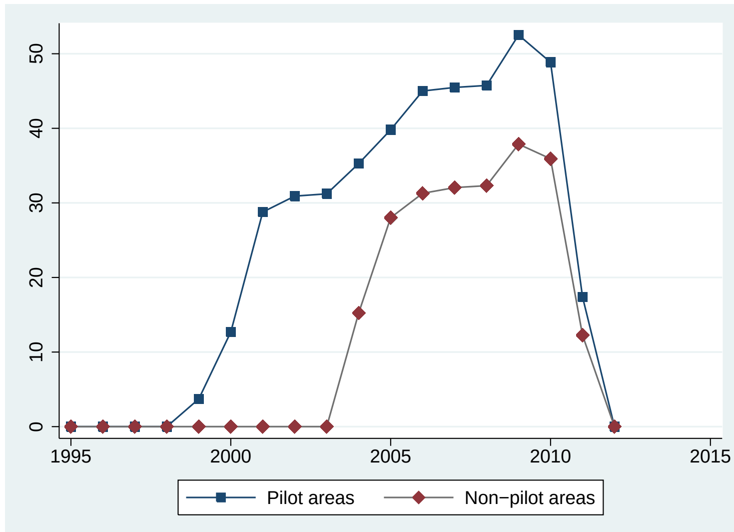

Notes: The graph displays the proportion of the area level population, which received EMA payments in the first two post-compulsory years of education, in EMA pilot and non-pilot areas by academic year.

Take-up substantially increased in the second pilot year, as these new students had prior knowledge that the programme was in place when making their school-leaving decisions, with only marginal increases seen in the subsequent pilot years. The national rollout saw further increases in take-up in the pilot areas, due to the increased generosity of the programme, with receipt increasing annually up to the academic year 2009/10, in non-pilot areas EMA receipt over the duration of the programme remained relatively constant in comparison. Take-up fell in the final years of the programme in England, as the allowance was closed to new applicants after the first quarter of the academic year 2010/2011.

A number of quantitative and qualitative evaluations of the impact of the programme in the initial pilot areas released each year of the pilot programme[6]. Using survey data collected during the pilot period Dearden, Emmerson, Frayne, and Meghir (2009) found that the impact of the programme was substantial, increasing full-time education rates by 4.5 and 6.7 percentage points in the first, and first two, years respectively. In analysis using administrative data Chowdry, Dearden, and Emmerson (2008), find smaller but still sizable impacts of the programme during the pilot period, a 3.0 (2.0) percentage point increase in the proportion of female (male) students. These estimates provide a useful benchmark for the area-level analysis in this study, see Table B.2 in Appendix B. Analysing the impact on juvenile crime rates in the first three years of the programme, Sabates and Feinstein (2007) find that the initial EMA pilot areas had a decrease of 1

[6]There were at least 16 different reports of the pilot evaluation, the final of which being Middleton, Perren, Maguire, Rennison, Battistin, Emmerson, and Fitzsimmons (2005).

conviction per 1,000 pupils in the burglary rate relative to non-EMA areas.

3.2.3 How could the EMA affect non-educational outcomes?

The EMA lowers the cost of post-compulsory education, and therefore we would expect an unambiguous increase in participation for the proportion of individuals eligible for the programme. Black et al. (2008) propose two mechanisms through which a legislative increase in the length of education can impact teenage fertility: the 'incarceration effect', where the requirement to spend longer at school reduces the opportunity to engage in risky behaviours; and the 'human capital effect', remaining in education raises an individual's stock of human capital, potentially improving their future economic prospects. In contrast to their setting, the EMA programme was not mandatory, so incarceration in this context has a different interpretation. For those individuals remaining in education the loss of opportunity for risky behaviour occurs as a result of a voluntary choice, thus by revealed preference it must be the case that the individual derives a higher utility from continued education than the forgone opportunity. Hence the direct effect of the EMA for these individuals is to increase the opportunity cost of non-educational activities. But although eligibility for the EMA was limited to only a sub-sample of the population, any increase in participation induced by the EMA would also induce a change to the peer group for not only the EMA participants, but also of individuals who in the absence of the programme would have continued education or who would have stopped schooling. Thus this indirect effect of the EMA to 'peer composition' potentially affects all individuals within the cohort. Hence the natural starting point of the analysis is to use the variation over time and areas in programme availability and generosity to investigate the impact of the programme on outcomes of the local population.

Two types of adolescent risk behaviour are considered in this study, teenage fertility and youth crime. Although teenage fertility rates may be affected by the behaviour of both adolescent males and female, it is an outcome that can be measured for females only. Conversely, youth crime is widely understood to be more prevalent amongst adolescent males. In the context of youth crime both the incarceration and human capital effect result in an increase in the opportunity cost of committing a crime. In addition to the change in peer composition effect outlined above, a longer schooling duration increases the interaction amongst young people with their peers, which may influence the proclivity towards criminal activity in either direction.

Similarly, the direction of the impact of an increase in post-compulsory participation on teenage fertility outcomes is unclear due to the change in student mix. Remaining in education also implies continued access to sexual health services at school or college facilities. Therefore those continuing education receive a potential reduction in contraception costs, which should have a downward effect on the conception rate. But as a pregnancy is a consequence of a social interaction, the change in peer group may result in a change in the frequency of social interactions, which could impact the conception rate in either direction. As the maternity rate is determined by the conception rate and the abortion decision, the impact on maternities is also ambiguous. However the prior on the percentage of conceptions leading to abortion is that there would be an unambiguous positive effect. The fertility decisions of young women who leave school at the compulsory leaving age would not be affected by the EMA. Those who are induced into post-compulsory education, who consequently become pregnant, are unlikely to be less likely to terminate a pregnancy than if they had not continued in education. For any individuals who were ineligible for the EMA, remained at school past compulsory age and became pregnant as a result of the change in peer group induced by the EMA, it is plausible that they would have a higher aversion to early motherhood, and hence be more likely to terminate a pregnancy.

3.3 Data

In order to evaluate the impact of the programme the ideal dataset would contain individual level measures of past and current EMA receipt and participation in post-compulsory education, as well fertility and crime measures. However this information is not available in one source, and therefore the analysis relies on two separate datasets. Individual-level responses regarding participation and fertility are compiled from the Labour Force Survey (LFS), which also contains information on household income and area of residence enabling the determination of whether an individual was eligbile for EMA receipt. Unfortunately there is no indicator in the LFS regarding whether an individual actually received the EMA. As data regarding the proportion of individuals receiving the EMA in each area is available a second dataset is compiled at the area level, using administrative data from England, Scotland and Wales at the top-tier Local Authority (LA) level from a range of official sources.

3.3.1 Individual-level Data

The dataset for the individual-level analysis combines data from the 1995-2013 Quarterly Labour Force Surveys. The survey is the largest representative survey in the UK, and contains detailed information of each individual within a household, including income, labour force status, gender, birth-date and ethnicity. Each household remains in the survey for five successive quarters, but as income information is available in the first and fifth wave only, individuals enumerated in these waves only are included in the sample[7]. The data is divided into sub-samples according to the academic birth cohort of the individual. Indicator variables are constructed using the geographical identifiers available within the secure data environment to indicate whether the area in which an individual lives had the EMA programme, and using information on the household income eligibility thresholds applicable to the area of residence to indicate whether an individual resides in a household eligible for EMA.

Table 3.1: Descriptive Statistics of the individual-level sample

	All		Males		Females	
	Mean	SD	Mean	SD	Mean	SD
FT education - year 1	0.758	0.429	0.718	0.450	0.799	0.400
FT education - year 2	0.637	0.481	0.598	0.490	0.678	0.467
Ethnicity:White	0.892	0.311	0.893	0.309	0.891	0.312
Ethnicity:Black	0.025	0.158	0.025	0.155	0.026	0.160
Ethnicity:Asian	0.051	0.219	0.049	0.216	0.053	0.223
Ethnicity:Other	0.032	0.176	0.033	0.180	0.030	0.172
British	0.952	0.215	0.953	0.212	0.950	0.217
EMA available	0.499	0.500	0.494	0.500	0.504	0.500
Household eligible for EMA	0.689	0.463	0.690	0.463	0.689	0.463
Household eligible for max EMA	0.469	0.499	0.466	0.499	0.472	0.499
Non-Working Household	0.199	0.399	0.196	0.397	0.201	0.401
Parent education - high	0.389	0.487	0.392	0.488	0.386	0.487
Mother					0.047	0.215
EMA available at 16					0.453	0.497
	N=45,636		N=23,243		N=22,393	

Descriptive statistics are displayed in Table 4.1 for individuals in the first two post-compulsory schooling years. The table reveals substantial age and gender heterogeneity in the proportion of individuals enrolled as a full-time student. In the first year after the compulsory education requirement 76% of the sample report being in education, which declines to 64% a year later. Females in both years are 8 percentage points more likely to be a student than males. In contrast there are no gender differences in individual characteristics such as ethnicity or nationality. Half the sample resides in a local authority which has implemented the EMA programme. Household characteristics are also similar across gender: approximately 70% of individuals in the sample reside in a household with income below the threshold requirement for programme eligibility, 47% in a household with income lower than the threshold for the maximum EMA payment. The proportion of individuals in households with no working parents is 20%, 39% of parents have quali-

[7]Through the longitudinal structure of the LFS an individual may be enumerated more than once in the sample. However as the estimation considers individuals in a single academic year, using the 1st and 5th waves of the survey implies that the individual is observed only once in any academic year.

fications that exceed the basic academic qualification level obtained through compulsory schooling. For the fertility analysis, a variable is constructed to indicate whether the EMA was available in the area of residence when the individual was in her first post-compulsory academic year. The proportion of mothers in the sample is 4.6%.

3.3.2 Area-level Data

For the area-level analysis a panel dataset is compiled of annual (academic year) area averages. Due to a large restructure of Local Government boundaries in April 1996, the data commences in academic year 1995/96 as the previous area structure is not directly comparable to the current Local Authority boundaries. Programme take-up, fertility, crime and post-compulsory participation rates are combined with area-level demographic variables from the LFS, aggregated using the supplied person weights.

Fertility measures were obtained from the Office of National Statistics (ONS), who release annual teenage conception rates, defined as a conception occurring to an individual aged less than 18, and the proportion of these conceptions that led to a recorded outcome either a maternity (live or stillbirth) or a legal abortion. Bespoke tabulations were commissioned for the analysis which recalculated the annual data to approximate the academic year[8]. The date and age at conception is calculated from the birth date in the case of a maternity, where a gestation time of 38 weeks is assumed. In the case of an abortion, date of birth of the woman and precise gestation of the foetus is recorded, from which the date and age at conception is determined. Miscarriages, and conceptions terminated through the use of emergency contraceptives (the day-after pill) are not recorded. However as long as the EMA did not induce a change in the miscarriage rate or the use of emergency contraceptives, the distribution of the outcome variable with regard to these measures will be unaffected, and the validity of the estimation remains. Since the Abortion Act (1967) the termination of a pregnancy is legal in the UK, and has been available via the centrally funded National Health Service since 1974. Therefore it is unlikely that individuals would resort to a non-registered provision. There are a small number of areas where one or more of the fertility measures is suppressed due to ONS confidentiality regulations, these are generally small areas with an event count of 10 occurrences or fewer.

The crime measure considered in the analysis is the rate of first time entrants to the criminal justice system, defined as individuals receiving their first reprimand, warning or conviction. The Ministry of Justice have collated this data from the Police National Computer (PNC) at a quarterly periodicity since 2000. The local authority data is calculated

[8] Teenage fertility measures are collected continuously, and collated at a quarterly periodicity. To approximate the academic year the bespoke tabulations combined the number of occurrences from the 4th quarter (October-December) in one calendar year with quarters 1-3 (January-September) in the following calendar year, using the mid-year population estimates for females aged 15-17 to calculate the rates.

using the home address of the offender as recorded by police, or the postcode of where the offence was committed if no home address was recorded. The data is available for two age groups: young people aged 10-17, and adults.

Participation in post-compulsory education rates were calculated using the cohort-specific number of participants, obtained from the Department of Education (DoE)[9] and ONS mid-year population estimates. Inner-London local authorities are excluded from the sample, as the DoE did not release this information prior to 2009. A further participation measure for each area in each period is available through aggregating the individual responses of current full-time student status from the LFS, where the aggregation uses the LFS supplied person weights.

The number of individuals in receipt of the EMA for each post-compulsory year and for each area during the post national rollout period was obtained from the Young People's Learning Agency (YPLA) and the DoE. For the pilot period, Chowdry et al. (2008) kindly shared the take-up data from academic years 2000/01-2003/04 used in their study. The remaining data was obtained after consulting the Hansard record of parliamentary debates and in response to requests from individual pilot local authorities. The proportion of households eligible for the EMA in each area was aggregated from the LFS using the person-income weights provided. Using this information and the ONS mid-year population estimates age-specific programme take-up rates, defined as the proportion of eligible individuals participation in the programme, were calculated.

Table 3.2: Descriptive Statistics of the area-level sample

	Mean	SD		Mean	SD
Educational Outcomes			**Non-Educational Outcomes**		
F/T education rate - 16	76.06	9.17	Conception rate	4.29	1.37
F/T education rate - 17	63.92	10.14	Maternity rate	2.33	0.97
Proportion F/T students - 16	76.49	18.02	% Abortions	47.80	9.74
Proportion F/T students - 17	64.63	21.22	First-time offenders : 10-17	1.43	0.67
			First-time offenders : adult	0.53	0.15
EMA measures			**Demographics**		
EMA available	56.56	49.58	Male: female ratio	1.05	0.11
Pilot area	37.84	48.51	Ethnicity: white	87.62	14.58
Proportion take-up: 16	39.41	11.37	British	87.84	12.24
Proportion take-up: 17	32.17	10.29	Real Hourly Wage: female 16-25	6.02	1.49
Proportion eligible	81.10	7.74	Real Hourly Wage: male 16-25	6.34	1.52
Proportion eligible max payment	64.99	8.29	Unemployment rate: female 16-25	14.19	8.44
			Unemployment rate: male 16-25	18.30	9.97

Table 3.2 displays descriptive statistics for the area-level sample. The full-time education rate in the first post-compulsory year (students aged 16 at the start of the academic

[9]The DoE data is a snapshot of the number of students in each academic cohort enumerated in December of each academic year.

year) is 76% and decreases to 64% in the second year after mandatory schooling. These administrative measures are comparable with the sample averages in the individual level sample in Table 4.1. The alternative area measures of participation in post-compulsory education, the proportion of full-time students which are calculated from the individual-level LFS dataset, provide a good approximation to the administrative data, although the data is less precise as evidenced by the larger standard deviations.

The fertility measures indicate that 4.3% of females under age 18 experience a pregnancy, 52% of which lead to a maternity. The proportion of first time entrants to the criminal justice system is 1.43% (0.53%) for individuals aged 10-17 (adults). There is also substantial between-area variation in the non-educational outcomes, see Figures B.2 and B.3 in Appendix B.

3.4 Empirical Methodology

As no single dataset measures fertility, participation and crime outcomes as well as EMA receipt, the analysis is comprised of two parts. Firstly an individual level analysis is performed to analyse the impact of being eligible for EMA receipt. In this strategy it is possible to evaluate the impact of the programme on post-compulsory participation in education as well as the propensity for early fertility. In the second part an area-level analysis facilitates the estimation of the impact of an increase in the post-compulsory education participation rate on area-level underage fertility and crime rates.

3.4.1 Individual-level Analysis

To estimate the influence of the EMA programme on the probability that an individual is a full-time student, a linear probability model is estimated, where the basic specification follows a difference-in-differences approach:

$$\begin{aligned} y_{ijt} = \beta_0 + \beta_1 EMA_{jt} + \beta_2 EligibleHH_{it} + \beta_3 EMAeligibleHH_{ijt} \\ + \gamma' X_{it} + \delta' Zjt + \alpha_j + \eta_t + \epsilon_{ijt} \end{aligned} \tag{3.1}$$

where y_{ijt} is a dummy variable indicating whether individual i living in area j at time t is currently enrolled as a full-time student; EMA_{jt} is an indicator variable taking the value 1 if the EMA programme was available in area j at time t and 0 otherwise; $EligibleHH_{it}$ indicates whether individual i at time t resides in a low income household (for periods prior to EMA implementation this variable is determined using the income threshold applicable to the area in the first year of the programme); the interaction term $EMAeligibleHH_{ijt}$ therefore takes on a value of 1 for individuals living in EMA eligible households in areas where the programme was available only; X_{it} and Z_{jt} are vectors of individual and area-level characteristics respectively; α_j and η_t are respectively area and time dummies; ϵ_{ijt} is an idiosyncratic error term. The coefficient of interest is therefore β_3, which should be interpreted as the Intention to Treat (ITT) effect, as it describes the impact of being eligible for the programme in an area and time period where the EMA was available

This basic specification is amended in order to investigate whether the programme had a differential impact by household income level by replacing the single eligibility indicator with two variables indicating whether the individual was eligible for the maximum weekly payment or for less than the maximum payment, and similarly replacing the interaction term:

$$\begin{aligned} y_{ijt} = \beta_0 &+ \beta_1 EMA_{jt} + \beta_2 MaxEligibleHH_{it} + \beta_3 MinEligibleHH_{it} \\ &+ \beta_4 MaxEMAeligibleHH_{ijt} + \beta_5 MinEMAeligibleHH_{ijt} \\ &+ \gamma' X_{it} + \delta' Zjt + \alpha_j + \eta_t + \epsilon_{ijt} \end{aligned} \tag{3.2}$$

where $MaxEligibleHH_{it}$ indicates whether household income falls below the lower EMA threshold, and $MinEligibleHH_{it}$ takes a value of 1 if household income falls between the upper and lower thresholds, and 0 otherwise. The associated interaction variables, $MaxEMAeligibleHH_{ijt}$ and $MinEMAeligibleHH_{ijt}$, therefore indicate whether an individual lives in the indicated household type when the programme is available. Thus in this specification the coefficients of interest are β_4 and β_5, which again describe the ITT effect for individuals eligible for the maximum and less than maximum weekly payments respectively.

Equations (3.1) and (3.2) allow the assessment of the influence of the programme on whether an individual is currently enrolled in full-time education. To investigate whether the EMA, through its impact on post-compulsory participation in education, influenced fertility behaviour is not straightforward, as the only observable measure of fertility is the presence of a child in the household, which can only be detected with a lag[10]. Therefore, to measure whether the programme exerted an influence on fertility, the estimation again proceeds via a linear probability model, with the basic specification amended to:

$$\begin{aligned} y_{ijt} = \beta_0 &+ \beta_1 EMAat16_{jt} + \beta_2 EligibleHH_{it} + \beta_3 EMAat16eligibleHH_{ijt} \\ &+ \gamma' X_{it} + \delta' Zjt + \alpha_j + \eta_t + \epsilon_{ijt} \end{aligned} \tag{3.3}$$

where $EMAat16_{jt}$ takes a value of 1 if the EMA was avaliable in the area in which an

[10] A change to behaviour inducing an effect on fertility which occurs at time t is not observable in the individual-level LFS data before $t + 9$ months.

individual resides when she was in her first post-compulsory academic year; as before $EligibleHH_{it}$ indicates a low income household; $EMAat16eligibleHH_{ijt}$ is the interaction term, taking a value of 1 if the EMA was available to the individual in the year after mandatory schooling and the individual resides in a low income household. This specification therefore relies on two assumptions, first that the individual currently resides in a household identical with regard to EMA eligibility as when she was 16 years of age, second that the current household financial situation is the same as when she was 16 years of age[11]. The coefficient of interest, β_3, describes the impact of the programme for those who were eligible to receive the EMA and therefore should again be interpreted as an Intention to Treat parameter.

Analogous to equation (3.2), this replacing the single eligibility indicator with two variables indicating eligibility by level of weekly EMA payment allows the investigation of whether the programme had a differential impact by household income level. The specification becomes:

$$\begin{aligned} y_{ijt} = \beta_0 &+ \beta_1 EMAat16_{jt} + \beta_2 MaxEligibleHH_{it} + \beta_3 MinEligibleHH_{it} \\ &+ \beta_4 MaxEMAat16eligibleHH_{ijt} + \beta_5 MinEMAat16eligibleHH_{ijt} \\ &+ \gamma' X_{it} + \delta' Zjt + \alpha_j + \eta_t + \epsilon_{ijt} \end{aligned} \tag{3.4}$$

so that the ITT effect are described by the coefficients on the relevant interaction terms, β_4 and β_5.

3.4.2 Area-level Analysis

The area level analysis investigates the impact of the EMA programme on post-compulsory education participation rates, first using variation in the availability of the EMA programme, second considering area-level variation in the intensity of the programme. The impact of changes in the post-compulsory participation rate, as induced by the EMA, on teenage fertility and crime rates is then considered.

3.4.2.1 Programme Availability Effect

The analysis first investigates the overall impact of the introduction of the EMA. As the programme implemented in the national rollout differed to that implemented during the pilot period, in terms of generosity of the payment structure and income threshold criteria,

[11] The individual and associated household income are observed in both wave 1 and wave 5 of the LFS. The data was therefore interrogated regarding the plausibility of these two assumptions, analysis using observations from wave 5 only combined with residence information from wave 1 produced consistent results.

the introduction of the pilot programme is not directly comparable to the introduction of the national rollout and therefore separate regressions are performed for pilot areas and non-pilot areas:

$$y_{jt} = \beta_0 + \beta_1 EMA_{jt} + g(TREND)_t + \delta' Z_{jt} + \alpha_j + \epsilon_{jt} \tag{3.5}$$

where y_{jt} measures the participation rate for area j at time t; EMA_{jt} is an indicator variable taking the value 1 if the EMA programme was available in area j at time t and 0 otherwise. As the EMA was introduced in areas contemporaneously, a polynomial function of an aggregate time trend, $g(TREND)$, is introduced to control for the national trend in participation; Z_{jt} is a vector of time-varying area-level characteristics; α_j control for time-invariant area characteristics; ϵ_{jt} is an area-specific error term. The coefficient of interest is β_1 which is identified by a jump in the education participation rate at the point of introduction, and cancellation, of the EMA programme in each area. This specification assumes that all areas experience a common time trend, but this assumption can be relaxed to allow for changing influences within an area over time[12], by the addition of a polynomial of region-specific time trends which captures local differences from the national trend. Estimation proceeds though Ordinary Least Squares (OLS).

Specification (3.5) identifies the impact of the EMA on participation solely through the introduction and cancellation of the programme, and therefore cannot determine whether the influence of the programme changed over time, for instance as a result of the national rollout of the programme. Hence an alternative specification is also considered to investigate the durational effects of the programme. In this approach all areas and time periods are pooled together and the single EMA binary indicator is replaced with a set of dummy variables denoting how many periods[13] the EMA was in place in a given area.

$$y_{jt} = \beta_0 + \sum_{k=1}^{M} \beta_k (EMA\ for\ k\ periods)_{jt} + \delta' Z_{jt} + \alpha_j + \eta_t + \epsilon_{jt} \tag{3.6}$$

As before the specification controls for time-varying area influences through a set of area fixed effects, α_j, and introduces a set of time dummies, η_t to flexibly account for aggregate effects common to all areas. As with specification (3.5) a polynomial of region-specific time trends can be included to account for local deviations from the aggregate time effects.

[12]Following Friedberg (1998), who used such a specification to investigate the impact of the introduction of unilateral divorce on state-level divorce rates

[13]This approach follows Wolfers (2003), who extended the analysis of Friedberg (1998) by using such a specification to investigated the durational effects associated with the introduction of unilateral divorce. Each period is defined as two academic years, coinciding with the maximum duration of EMA receipt for an eligible individual.

3.4.2.2 Programme Intensity Effect

The latter specification allows the analysis of the influence of the programme beyond just its introduction. The key difference between the programme implemented during the pilot period and that implemented in the national rollout was the eligibility criteria. Specifically the income threshold for the maximum payment was increased in order to encourage higher programme take-up. Furthermore, as described in Section 3.2.2, as the income thresholds for eligibility were kept in nominal terms further variation over areas and time in the proportion of the population eligible for the programme is induced through inflation. This area-level variation is exploited to measure of the intensity of eligibility on participation rates. The basic specification follows the difference-in-difference approach used in the individual level analysis:

$$y_{jt} = \beta_0 + \beta_1 EMA_{jt} + \beta_2 propbelow_{jt} + \beta_3 EMAbelow_{jt} + \delta' Z_{jt} + \alpha_j + \eta_t + \epsilon_{jt} \quad (3.7)$$

where y_{jt} measures the education participation rate in area j at time t; EMA_{jt} is an indicator variable taking the value 1 if the EMA programme was available in area j at time t and 0 otherwise; $propbelow_{jt}$ is the proportion of households in area j at time t with income level under the EMA eligibility threshold[14]; the interaction term $EMAbelow_{jt}$ therefore takes on positive values in area/time cells where the EMA was available, 0 otherwise; Z_{jt} is a vector of time-varying area-level characteristics; α_j and η_t are respectively area and time dummies; ϵ_{ijt} is an idiosyncratic error term. Region-specific time trends can be introduced to control for changing influences within a region over time. The coefficient of interest is this specification is β_3, which describes the Intention to Treat (ITT) effect, defined as the impact of the programme on the eligible population in an area during the period when the programme was available.

Administrative data is available regarding the actual number of recipients of the EMA payments by local authority and academic year. However, as the programme was voluntary, there may influences at the individual level that impact both the decision to remain in education and whether they apply for the programme or not which are not observable at the area level, implying that an OLS regression of the education participation rate on the programme participation rate would be biased. Therefore a two stage least squares (2SLS) approach is taken, whereby the actual take-up rate in an area, defined as the number of recipients of the EMA as a proportion of the *eligible* population, is instrumented with exogenous changes in the proportion eligible for EMA receipt. The two-step approach can

[14] As with the individual-level analysis, for periods prior to EMA implementation this variable is determined using the income threshold applicable to the area in the first year of the programme.

be written as:

$$Takeup_{jt} = \alpha_0 + \alpha_1 EMA_{jt} + \alpha_2 propbelow_{jt} + \alpha_3 EMAbelow_{jt} + \delta' Z_{jt} + \alpha_j + \eta_t + \epsilon_{jt} \quad (3.8)$$

$$EducRate_{jt} = \xi_0 + \xi_1 \widehat{Takeup}_{jt} + \phi' Z_{jt} + \tau_j + \rho_t + e_{jt} \quad (3.9)$$

where in the first stage, (3.8), the impact of changes in the eligible proportion is estimated, and then included in the second stage equation (3.9). Therefore ξ_1 describes the causal effect of an increase in take-up of the EMA on post-compulsory participation.

3.4.2.3 Non-educational outcomes

There may be unobservable factors that affect not only the non-educational outcomes of interest, fertility and crime rates, but also influence post-compulsory education rates. To address this potential endogeneity, the estimation of the effect of participation on these non-educational outcomes proceeds via 2SLS. In the first stage the impact of exogenous changes in the proportion eligible for the EMA on the education participation rate is estimated following specification (3.7). The second stage becomes:

$$y_{jt} = \psi_0 + \psi_1 \widehat{EducRate}_{jt} + \delta' Z_{jt} + \alpha_j + \eta_t + \epsilon_{jt} \quad (3.10)$$

where y_{jt} measures the outcome rate of interest in area j at time t; Z_{jt}, α_j, η_t and ϵ_{jt} are defined as in (3.7). Thus ψ_1 captures the effect of an increase in the post-compulsory education participation rate on the outcomes of interest.

3.5 Results

Section 3.5.1 presents the results of the individual-level analysis, where the intention to treat effect, defined as the average response among the would-be eligible population in areas and periods where the EMA was available, on whether an individual reports currently being a full-time student, and motherhood is estimated. The area-level analysis is presented in Section 3.5.2, presenting first the analysis of the impact of the availability of the EMA programme and area-level intensity of eligibility for the programme on post-compulsory participation in education. Finally the analysis explores the impact of changes in area-level post-compulsory education rates induced by the EMA programme on non-educational outcomes.

3.5.1 Individual-level Analysis

Table 3.3 presents the results for individuals observed in the first year after compulsory education. The upper panel reports the estimates for all individuals pooled together, whereas the middle and bottom panels report the estimates for males and females respectively. Columns (1)-(4) analyse the overall impact of being eligible for the EMA, whereas columns (5)-(8) report separate estimates according to whether the individual is eligible for the maximum weekly EMA payment or for the lower EMA payments. All specifications control for area and time fixed effects, columns (2)-(4) and (6)-(8) include individual level controls (nationality, ethnicity, gender and parental education). Columns (3)-(4) and (7)-(8) also include time-varying area controls (the gender ratio, gender-specific unemployment rates and real hourly wage (both for the 16-25 age group), area-level ethnicity and nationality). Specifications (4) and (8) also include area-specific linear trends, and hence produce the most conservative results.

The estimates in the upper panel reveal that an individual residing in a household with income below the EMA threshold, in an area where the programme was available, has a 3.3 percentage point higher likelihood of currently being enrolled as a full-time student, with a larger impact for individuals eligible for the maximum payments, although the difference is not statistically significant. As the estimates measure the impact of eligibility for the EMA programme rather than programme participation, they should be interpreted as the Intention to Treat Effect (ITT). Examining the ITT effects separately by gender reveals substantial gender heterogeneity, insofar with the impact for males is 4.9 percentage points, whereas the corresponding estimate for females is smaller and not significantly different from zero. Examining the impact of the programme by payment type reveals that the impact of the programme is larger, 6.5 percentage points, for males from the lowest income households.

Table 3.4 presents the analogous estimates for individuals in the second post-compulsory academic year. The overall impact in the second year is larger than in year 1, with eligibility associated with a 5.7 percentage point increase in the probability of being in full-time education, and again a slightly larger effect for individuals eligible for the maximum payment. In contrast to the year 1 findings presented in Table 3.3, the response in the second post-compulsory reveals little gender heterogeneity in the overall impact of eligibility for the programme. For males there is no significant difference in the response by income threshold, whereas the female ITT effect is driven almost entirely by the response of individuals available for the maximum payment.

A comparison of the estimates over the two years indicates that for males the impact of the programme was approximately equal in each year, whereas the programme elicited a significant effect for females only in the second post-compulsory year. That the increase

Table 3.3: Enrolled as full-time student: post-compulsory year 1

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Eligibile*EMA	0.035***	0.035***	0.035***	0.033***				
	(0.011)	(0.011)	(0.011)	(0.011)				
MaxEligible*EMA					0.043***	0.042***	0.042***	0.039***
					(0.012)	(0.012)	(0.012)	(0.012)
MinEligible*EMA					0.038**	0.032**	0.032**	0.032**
					(0.015)	(0.015)	(0.015)	(0.015)
N	23,611	23,611	23,611	23,611	23,611	23,611	23,611	23,611
Males								
Eligibile*EMA	0.059***	0.053***	0.053***	0.049***				
	(0.016)	(0.016)	(0.016)	(0.016)				
MaxEligible*EMA					0.074***	0.067***	0.068***	0.065***
					(0.018)	(0.017)	(0.017)	(0.018)
MinEligible*EMA					0.058**	0.045**	0.044**	0.041*
					(0.023)	(0.022)	(0.022)	(0.022)
N	12,060	12,060	12,060	12,060	12,060	12,060	12,060	12,060
Females								
Eligibile*EMA	0.008	0.010	0.010	0.010				
	(0.015)	(0.014)	(0.015)	(0.015)				
MaxEligible*EMA					0.012	0.013	0.013	0.011
					(0.016)	(0.016)	(0.016)	(0.016)
MinEligible*EMA					0.009	0.009	0.009	0.012
					(0.021)	(0.021)	(0.021)	(0.021)
N	11,551	11,551	11,551	11,551	11,551	11,551	11,551	11,551
Individual controls	no	yes	yes	yes	no	yes	yes	yes
Area controls	no	no	yes	yes	no	no	yes	yes
Area trends	no	no	no	yes	no	no	no	yes

Notes: Columns (1)-(4) display estimates of equation (3.1), whereas columns (5)-(8) presents the results from equation(3.2), as described in Section 3.4.1. All specifications include area and time fixed effects. The vector of individual controls comprizes nationality, ethnicity, gender and parental education; area controls include the gender ratio, gender specific unemployment rates and real hourly wage, as well as area-level UK nationality and ethnic population proportions. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, ***$p < 0.01$

is larger for males in the first year of the allowance may reflect the lower staying-on rate for males relative to females as indicated in Table 4.1. In contrast, eligibility for the EMA appears to have had a delayed effect on the likelihood of current enrollment for females, which is more likely driven by a reduction in dropout rather than new entrants to post-compulsory education. That the estimates are largest for individuals eligible for the maximum weekly payment suggests that the EMA was particularly successful at encouraging an increase in participation of individuals from lower income households.

Table 3.5 reports the results of the analysis on teenage fertility. The sample is restricted to females only, observed in the second and third post-compulsory years (aged 17 or 18 at the beginning of the relevant academic year), as the measured outcome, whether an

Table 3.4: Enrolled as full-time student: post-compulsory year 2

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Eligibile*EMA	0.067***	0.062***	0.062***	0.057***				
	(0.013)	(0.013)	(0.013)	(0.013)				
MaxEligible*EMA					0.076***	0.071***	0.071***	0.065***
					(0.014)	(0.014)	(0.014)	(0.014)
MinEligible*EMA					0.066***	0.053***	0.052***	0.049***
					(0.018)	(0.017)	(0.017)	(0.018)
N	22,025	22,025	22,025	22,025	22,025	22,025	22,025	22,025
Males								
Eligibile*EMA	0.066***	0.060***	0.060***	0.055***				
	(0.019)	(0.019)	(0.019)	(0.019)				
MaxEligible*EMA					0.067***	0.060***	0.060***	0.055**
					(0.021)	(0.021)	(0.021)	(0.021)
MinEligible*EMA					0.076***	0.062***	0.061**	0.058**
					(0.026)	(0.025)	(0.025)	(0.026)
N	11,183	11,183	11,183	11,183	11,183	11,183	11,183	11,183
Females								
Eligibile*EMA	0.061***	0.059***	0.059***	0.052***				
	(0.019)	(0.018)	(0.018)	(0.019)				
MaxEligible*EMA					0.082***	0.079***	0.079***	0.071***
					(0.020)	(0.020)	(0.020)	(0.020)
MinEligible*EMA					0.038	0.032	0.032	0.027
					(0.025)	(0.024)	(0.024)	(0.024)
N	10,842	10,842	10,842	10,842	10,842	10,842	10,842	10,842
Individual controls	no	yes	yes	yes	no	yes	yes	yes
Area controls	no	no	yes	yes	no	no	yes	yes
Area trends	no	no	no	yes	no	no	no	yes

Notes: see notes to Table 3.3

individual is a mother, is observable with a delay[15]. The results indicate that those residing in households eligible for the EMA in areas where the EMA would have been available in the individual's first post-compulsory year are 0.9 percentage points less likely to be a mother. This represents a decrease of approximately 19% relative to the sample mean, with the reduction driven specifically by those individuals eligible for the higher EMA payment. Surprisingly a positive estimate is found for individuals eligible for the lower weekly EMA payment, although this estimate is not significantly different from zero.

[15] If the sample is restricted to individuals in the second post-compulsory year only the estimates are consistent with those presented here, although the magnitudes become slightly smaller and lose precision due to the smaller sample size.

Table 3.5: Impact of the EMA programme on individual-level fertility

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Eligibile*EMAat16	-0.010*	-0.010*	-0.009*	-0.009*				
	(0.006)	(0.006)	(0.006)	(0.006)				
MaxEligible*EMAat16					-0.037***	-0.037***	-0.037***	-0.036***
					(0.007)	(0.007)	(0.007)	(0.007)
MinEligible*EMAat16					0.009	0.008	0.008	0.007
					(0.006)	(0.007)	(0.007)	(0.007)
N	20,047	20,047	20,047	20,047	20,047	20,047	20,047	20,047
Individual controls	no	yes	yes	yes	no	yes	yes	yes
Area controls	no	no	yes	yes	no	no	yes	yes
Area trends	no	no	no	yes	no	no	no	yes

Notes: Columns (1)-(4) display estimates of equation (3.3), whereas columns (5)-(8) presents the results from equation(3.4), as described in Section 3.4.1. All specifications include area and time fixed effects. The vector of individual controls comprizes nationality, ethnicity, gender and parental education; area controls include gender specific unemployment rates and real hourly wage, as well as area-level UK nationality and ethnic population proportions. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, ***$p < 0.01$

3.5.2 Area-level Analysis

The analysis now turns to the area-level analysis. The results explore first the effect of the programme's availability, and second the intensity of the programme on post-compulsory education participation rates. Finally, the impact of exogenous variation in participation induced by the EMA programme on non-educational outcomes is explored.

Table 3.6 presents impact of the availability of the EMA programme on area-level full-time education participation rates. Panel A considers the impact of the introduction of the EMA programme, where the analysis is performed separately for pilot and non-pilot areas, as the programme implemented in pilot areas differed from that implemented in non-pilot areas in terms of eligbility criteria (see Section 3.2.2). Panel B presents the durational effects of the programme, where the estimation pools pilot and non-pilot areas and thus features variation in the timing of the introduction and cancellation of the EMA programme. Columns (1)-(3) refer to the first post-compulsory year, (4)-(6) to the second post-compulsory year, and (7)-(9) to the overall impact on both the first and second post-compulsory year. All specifications include a set of individual area fixed effects; Panel A includes a quadratic polynomial in the aggregate time trend[16], whereas Panel B controls for aggregate time influences by including a set of year fixed effects. The second and third columns within each set of results include a vector of area level controls (the male to female ratio, the proportion of the population with British nationality, ethnic proportions, male and female unemployment rates and real hourly wages for the 16-25 age group). The third column also includes a set of region-specific linear time trends. All regressions are

[16]The appropriate form of the polynomial was determined using the Akaike Information Criterion. The results are however also robust to using a linear or cubic polynomial.

Table 3.6: Programme Availability Effect

	Year 1 Participation			Year 2 Participation			Year 1&2 Participation		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Program Availability									
Non-Pilot	3.930***	3.392***	3.449***	2.731***	2.487***	2.565***	2.869***	2.288***	2.372***
	(0.264)	(0.249)	(0.246)	(0.255)	(0.266)	(0.263)	(0.233)	(0.241)	(0.239)
N	1432	1426	1426	1436	1430	1430	1433	1427	1427
Pilot areas	4.138***	4.894***	4.853***	2.430***	2.000***	2.095***	3.011***	2.122***	2.189***
	(0.751)	(0.703)	(0.676)	(0.602)	(0.611)	(0.567)	(0.558)	(0.568)	(0.520)
N	753	751	751	753	751	751	753	751	751
Panel B: Program Duration									
Period 1	0.486	0.808*	1.339***	-0.302	-0.066	0.400	0.040	0.308	0.804**
	(0.460)	(0.435)	(0.423)	(0.451)	(0.437)	(0.423)	(0.409)	(0.390)	(0.374)
Period 2	2.897***	2.211***	2.719***	1.521***	1.431***	1.886***	2.153***	1.795***	2.276***
	(0.502)	(0.468)	(0.450)	(0.437)	(0.420)	(0.391)	(0.421)	(0.395)	(0.368)
Period 3	4.425***	4.013***	4.394***	3.566***	3.613***	3.957***	3.954***	3.794***	4.155***
	(0.695)	(0.650)	(0.619)	(0.643)	(0.621)	(0.579)	(0.600)	(0.568)	(0.526)
Period 4	7.779***	6.753***	6.231***	6.523***	6.188***	5.731***	7.101***	6.461***	5.973***
	(0.963)	(0.900)	(0.844)	(0.921)	(0.890)	(0.828)	(0.842)	(0.799)	(0.734)
N	2,185	2,177	2,177	2,189	2,181	2,181	2,186	2,178	2,178
Area controls	no	yes	yes	no	yes	yes	no	yes	yes
Region trends	no	no	yes	no	no	yes	no	no	yes

Notes: Panel A reports estimation of equation (3.5), Panel B uses equation (3.6) both as described in Section 3.4.2.1. All specifications include area fixed effects and a quadratic aggregate year trend, and are weighted according to area population size. The vector of area controls includes the gender ratio, gender specific unemployment rates and real hourly wage, as well as area-level UK nationality and ethnic population proportions. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, ***$p < 0.01$

weighted with the size of the age 16-17 population.

The results indicate that for the first year of post-compulsory education the introduction of the EMA is associated with a larger increase in the education rate in pilot areas relative to non-pilot areas, 4.9 and 3.5 percentage points respectively, which may reflect that pilot areas had lower initial post-compulsory education rates. In contrast to the individual analysis results, the area analysis suggests that the impact of the EMA was lower in the second year of the programme, estimated as 2.6 (2.1) percentage points for pilot (non-pilot) areas, although the difference is not statistically significant. If the analysis is restricted to the pilot period only, see Table B.2 in Appendix B, the estimates of the effect of the introduction of the programme on pilot areas are consistent with those presented for non-pilot areas in Table 3.6. Whereas the estimates in Panel A describe the average impact for each year of the programme, the duration effect presented in Panel B illustrates the effect of the EMA programme over consecutive cohorts. These latter indicate that the impact of the EMA on post-compulsory education participation rates was negligible in the first years of the policy, but increased with the duration of the programme. However as the national rollout occurred in period 3 for pilot areas, whereas for non-pilot areas the national rollout happened in period 1, it is not clear whether the increased impact of the EMA is driven by the longer availability of the programme or by the significant changes to the eligibility criteria associated with the national rollout, which imply that a greater proportion of individuals were eligible for the maximum payment. This latter can be understood as an increase in the intensity of the programme. The area analysis therefore turns to examining the effect of the intensity of the EMA within areas, using the proportion of the population eligible for the EMA as the policy variable.

Table 3.7: Programme Intensity Effect

	Year 1 Participation				Year 2 Participation				Year 1&2 Participation			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: F/T Education												
EMA eligibility	0.261***	0.224***	0.164***	0.175***	0.202***	0.210***	0.137***	0.146***	0.228***	0.214***	0.130***	0.135***
	(0.031)	(0.029)	(0.028)	(0.029)	(0.028)	(0.027)	(0.028)	(0.028)	(0.026)	(0.025)	(0.025)	(0.025)
N	2,185	2,177	2,177	2,177	2,189	2,181	2,181	2,181	2,186	2,178	2,178	2,178
EMA Take-up	0.219***	0.188***	0.140***	0.162***	0.215***	0.219***	0.152***	0.167***	0.249***	0.233***	0.151***	0.171***
	(0.025)	(0.025)	(0.024)	(0.026)	(0.029)	(0.028)	(0.032)	(0.032)	(0.030)	(0.030)	(0.030)	(0.032)
N	2,159	2,151	2,151	2,151	2,165	2,157	2,157	2,157	2,027	2,021	2,021	2,021
Panel B: Proportion F/T students												
EMA eligibility	0.154*	0.153*	0.103	0.063	0.193*	0.162	0.109	0.111	0.158**	0.148**	0.082	0.070
	(0.083)	(0.085)	(0.087)	(0.092)	(0.104)	(0.106)	(0.116)	(0.116)	(0.065)	(0.067)	(0.072)	(0.072)
N	2,475	2,462	2,462	2,462	2,479	2,464	2,464	2,464	2,487	2,471	2,471	2,471
EMA Take-up	0.139*	0.138*	0.093	0.061	0.222*	0.172	0.124	0.122	0.213**	0.183**	0.107	0.106
	(0.077)	(0.078)	(0.083)	(0.090)	(0.123)	(0.120)	(0.142)	(0.143)	(0.084)	(0.083)	(0.093)	(0.099)
N	2,446	2,433	2,433	2,433	2,452	2,437	2,437	2,437	2,312	2,300	2,300	2,300
Panel C: Proportion F/T male students												
EMA eligibility	0.158	0.146	0.090	0.053	0.131	0.080	0.025	0.027	0.164*	0.141**	0.065	0.058
	(0.133)	(0.136)	(0.140)	(0.145)	(0.152)	(0.156)	(0.169)	(0.171)	(0.096)	(0.098)	(0.103)	(0.105)
N	2,417	2,405	2,405	2,405	2,394	2,383	2,383	2,383	2,479	2,464	2,464	2,464
EMA Take-up	0.142	0.130	0.071	0.049	0.153	0.082	0.029	0.030	0.246**	0.201*	0.117	0.132
	(0.122)	(0.125)	(0.134)	(0.142)	(0.178)	(0.173)	(0.205)	(0.208)	(0.122)	(0.120)	(0.133)	(0.142)
N	2,390	2,378	2,378	2,378	2,367	2,356	2,356	2,356	2,304	2,293	2,293	2,293
Panel D: Proportion F/T female students												
EMA eligibility	0.153	0.140	0.119	0.096	0.225	0.216	0.156	0.152	0.181**	0.181*	0.141	0.121
	(0.121)	(0.126)	(0.134)	(0.141)	(0.143)	(0.146)	(0.157)	(0.159)	(0.090)	(0.093)	(0.100)	(0.101)
N	2,405	2,394	2,394	2,394	2,399	2,387	2,387	2,387	2,475	2,461	2,461	2,461
EMA Take-up	0.138	0.124	0.107	0.090	0.252	0.232	0.175	0.164	0.225*	0.209*	0.153	0.140
	(0.111)	(0.116)	(0.127)	(0.138)	(0.167)	(0.162)	(0.189)	(0.193)	(0.116)	(0.114)	(0.128)	(0.136)
N	2,376	2,365	2,365	2,365	2,372	2,360	2,360	2,360	2,302	2,291	2,291	2,291
Area controls	no	yes	yes	yes	no	yes	yes	yes	no	yes	yes	yes
Linear trends	no	no	yes	yes	no	no	yes	yes	no	no	yes	yes
Quadratic trends	no	no	no	yes	no	no	no	yes	no	no	no	yes

Notes: The table reports estimation of equations (3.7) and (3.9), for EMA eligibility and EMA Take-up respectively, as described in Section 3.4.2.2. All specifications include area and year fixed effects, and are weighted according to area population size. The vector of area controls includes the gender ratio, gender specific unemployment rates and real hourly wage, as well as area-level UK nationality and ethnic population proportions. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, ***$p < 0.01$

The results of the analysis of the area-level programme intensity are displayed in Table 3.7. Analogous to the individual-level analysis this measures the Intention to Treat (ITT) effect, as it considers the impact of changes in the proportion of individuals eligible for the treatment rather than those who received the treatment. As there is both spatial and temporal variation in eligibility, the specification mirrors that used in the individual analysis. Administrative data is also available regarding the number of individuals in each academic cohort who actually received the EMA in each area, allowing the calculation of a measure of actual programme take-up (the proportion of eligible individuals in receipt of the EMA). However as there may be unobserved influences that impact both the decision to apply for the EMA as well as the participation decision, a two-stage least squares (2SLS) approach is taken whereby the actual take up rate is instrumented with exogenous changes in the proportion eligible for the programme. The first specification in each set of results includes just area and year fixed effects. The second specification adds the vector of time-varying area level controls. The third and fourth specification add individual linear and quadratic region-specific time trends respectively.

Considering the impact on the full-time education participation rate in Panel A, the estimates are positive and significant and slightly decrease in magnitude when region-specific time trends are included. In the first (second) post-compulsory year a one percentage point increase in the proportion of the population eligible for the EMA in areas where programme was present is associated with a 0.18 (0.15) percentage point increase in the age-specific education rate. The corresponding 2SLS estimates indicate that a one percentage point increase in programme take-up, induced by exogenous changes in eligibility, is associated with a 0.16 (0.17) percentage point increase in the appropriate education rate. That the OLS and 2SLS estimates are so similar is indicative that changes in actual take-up of the EMA programme can largely be explained by changes in eligibility criteria alone.

Panels B, C and D consider the proportion of full-time students, male students and female students. These measures are aggregated from the individual level Quarterly Labour Force Surveys, and are therefore not as precise as the administrative measure used in panel A. Nevertheless, the estimates for all students, panel B, are qualitatively consistent with those obtained with the administrative data. Examining the impacts for male and female students separately suggests that the changes in the eligibility criteria induced a larger response amongst females. Although less precise, these latter measures are available for all local authorities areas, whereas the administrative measure of the full-time education rate does not include the 11 Inner-London local authorities. The LFS measures are however useful as they allow for the assessment of gender-heterogeneity in the impact of the programme.

The programme intensity estimates presented in Table 3.7 become the first-stage estimates in the analysis of the influence of participation in post-compulsory education on teenage fertility and crime. As these non-educational outcomes are rates defined over the female population aged 15-17 for fertility measures and over the 10-17 youth population for crime, the impact of the EMA on participation for both the first and second post-compulsory years combined are considered as the first stage. Table 3.8 presents the results for three different fertility measures: the under-age (under 18) conception rate, maternity rate and the percentage of underage conceptions that led to abortion. The im-

Table 3.8: Impact of the EMA programme on area level fertility

	Conceptions				Maternities				% Led to Abortion			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
F/T Education	-0.033***	-0.047***	-0.030	-0.035	-0.065***	-0.075***	-0.070***	-0.073***	0.302***	0.337**	0.498*	0.509*
	(0.012)	(0.013)	(0.024)	(0.024)	(0.010)	(0.011)	(0.020)	(0.020)	(0.122)	(0.132)	(0.263)	(0.269)
N	2044	2040	2040	2040	2044	2040	2040	2040	2045	2041	2041	2041
Female Students	-0.032	-0.062*	-0.063	-0.064	-0.081*	-0.095*	-0.082	-0.083	0.400	0.428	0.502	0.518
	(0.022)	(0.035)	(0.046)	(0.047)	(0.042)	(0.049)	(0.056)	(0.058)	(0.257)	(0.270)	(0.402)	(0.518)
N	2325	2316	2316	2316	2325	2316	2316	2316	2326	2317	2317	2317
Area controls	no	yes	yes	yes	no	yes	yes	yes	no	yes	yes	yes
Linear trends	no	no	yes	yes	no	no	yes	yes	no	no	yes	yes
Quadratic trends	no	no	no	yes	no	no	no	yes	no	no	no	yes

Notes: The table reports estimation of equation (3.10) as described in Section 3.4.2.3. All specifications include area and year fixed effects, and are weighted according to area population size. The vector of area controls includes the gender ratio, gender specific unemployment rates and real hourly wage, as well as area-level UK nationality and ethnic population proportions. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, ***$p < 0.01$

pact of total participation on the teenage conception rate is negative in all specifications, but does not reach statistical significance, whereas the impacts on the maternity rate and the percentage of conceptions terminated are larger, and significantly different from zero. The estimates indicate that a one percentage point increase in the full-time education rate of individuals in the first and second years of post-compulsory schooling, induced by changes in the proportion of the population eligible for the EMA programme, is associated with a 0.07 percentage point decrease in the maternity rate, and a 0.51 percentage point increase in the proportion of conceptions leading to an abortion. As the full-time education participation rate does not discriminate between males and females, whereas the fertility measures are defined over females only, it may be more illuminating to examine the fertility response with respect to the proportion of female students only. Recall that this measure is aggregated from the LFS, and the resulting estimates are therefore less precise than those obtained with the administrative data. Nevertheless the estimates are qualitatively consistent, and indicate that a one percentage point increase in the proportion of females enrolled in full-time education is associated with a 0.06 (0.08) percentage point decrease in the conception (maternity) rate and a 0.52 percentage point increase in terminations. At the sample mean of eligibility these estimates translate to approximately a 19% decrease in the under 18 maternity rate, driven by both a decrease in conceptions (8%) and an increase in abortions (6%)[17].

The final analysis considers the influence of an increase in post-compulsory education on youth crime. The estimation follows the specification used in the fertility analysis, with the results displayed in Table 3.9. The estimates consistently indicate that the increase in participation is associated with a downward impact on the proportion of first-time youth offenders. However the estimates are not statistically significant. The impact of an

[17]These back of the envelope calculations use the first-stage estimate in column 12 of Table 3.7 and the proportion eligible for the EMA at the sample mean from Table 3.2

Table 3.9: Impact of the EMA programme on area level offenders

	Youth Offenders				Adult Offenders			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
F/T Education	-0.025	-0.027	-0.001	-0.019	-0.008*	-0.004	0.010	0.005
	(0.016)	(0.019)	(0.024)	(0.021)	(0.005)	(0.005)	(0.007)	(0.005)
N	1862	1854	1854	1854	1862	1854	1854	1854
F/T Students	-0.054	-0.053	-0.010	-0.023	-0.019	-0.006	0.012	0.003
	(0.095)	(0.069)	(0.036)	(0.028)	(0.033)	(0.012)	(0.018)	(0.006)
N	2043	2027	2027	2027	2043	2027	2027	2027
Male Students	-0.064	-0.081	-0.026	-0.054	-0.029	-0.010	0.034	0.007
	(0.179)	(0.213)	(0.173)	(0.153)	(0.063)	(0.029)	(0.193)	(0.022)
N	2037	2022	2022	2022	2037	2022	2022	2022
Female Students	-0.025	-0.027	-0.002	-0.011	-0.009	-0.003	0.005	0.002
	(0.034)	(0.029)	(0.013)	(0.013)	(0.012)	(0.006)	(0.005)	(0.003)
N	2032	2018	2018	2018	2032	2018	2018	2018
Area controls	no	yes	yes	yes	no	yes	yes	yes
Linear trends	no	no	yes	yes	no	no	yes	yes
Quadratic trends	no	no	no	yes	no	no	no	yes

Notes: See notes to Table 3.8

increase in the proportion of male students is estimated to be larger, consistent with a hypothesis that youth offenders are disproportionately male, but the estimates still do not achieve significance. As a robustness exercise the analysis is repeated for adult offenders, the resulting estimates are not statistically different to zero, but in contrast to the youth offending rates are also of a far smaller magnitude and not consistent in sign over the specifications. Although suggestive, there are some caveats to this final analysis, as the crime measure considered in the proportion of first-time entrants to the criminal justice system, in other words, the number of new offenders known to the police in each area and time period, rather than the incidence of crimes perpetrated. Additionally, as the rate is defined over a wide age range, it incorporates young people within compulsory education as well as those beyond the compulsory school-leaving age.

3.6 Conclusion

This paper has investigated the impact of the Education Maintenance Allowance (EMA) on participation rates in post-compulsory education and non-educational outcomes. Consistent with evaluations of the EMA pilots, the programme induced a substantial increase in the staying-on rates, especially amongst individuals eligible for the maximum EMA payments. There is also evidence of substantial gender heterogeneity, in the first-compulsory year the increase in EMA was particularly effective at increasing the participation rates amongst males, whereas in the second post-compulsory year the impact was larger for females. The analysis also considered the impact of post-compulsory participation on non-educational measures, teenage fertility and the proportion of first-time offenders. The results indicate that the increase in participation induced by changes in EMA eligibility is associated with a substantial decrease in the underage maternity rate, which is driven by both a reduction in the conception rate as well as an increase in the proportion of teenage conceptions that led to abortion. Weaker, but suggestive evidence is found that the programme influenced youth crime.

Part II

Marriage and Partnership

Chapter 4

Marriage Market Consequences of an Educational Reform

4.1 Introduction

Who to marry is one of the most important decisions a person undertakes, a decision that influences not only the individual's lifetime utility, but which also has important intergenerational consequences through the potential impact on the outcomes of children. Indeed, household formation can be understood as a key determinant of intergenerational mobility and income inequality: Greenwood, Guner, Kocharkov, and Santos (2014) show that the US Gini coefficient for household income would be significantly smaller if spouses matched randomly rather than assortatively. The quality of a specific match depends on an individual's relative position in the marriage market, which in turn is affected by factors such as family background and social standing, previous educational investments and future economic opportunities. Although these are idiosyncratic influences, two noticeable features emerge in the majority of marriages.

First, husbands are on average older than their wives. In all 218 countries for which data exists we observe a positive marital age gap[1]. Over the last 40 years the marital age gap has exhibited a downward trend in many countries, coinciding with strong increases in the mean age at marriage, however this positive age gap has endured. Indeed, even in countries with the highest standards of gender equality, such as the Scandinavian nations, an age gap of about two years persists. The second noticeable feature is that in general spouses have similar education levels. Schwartz and Mare (2005), with US data, show that the degree of positive assortative matching in education between spouses has increased since 1960, driven predominantly by increased homogamy in both tails of the marriage

[1] See Appendix C.1, Figure C.1

distribution: individuals with low educational attainment are more likely to have a low educated spouse, and similarly college graduates are increasingly likely to marry each other. Most striking is the decline in the likelihood that individuals with very low education levels form a marriage with a more educated spouse.

These two stylised facts imply an overlooked consequence for cohort specific educational reforms. As individuals typically match with partners from different cohorts, both spouses are not subject to the same institutional setting if a reform falls in between their birth dates. In turn, standard matches with positive age gaps and equality in education become infeasible. Who marries whom then depends on the relative importance of preferences over a prospective partner's age and homogamy in education. This may have non-trivial consequences for reform evaluation, as direct benefits of the reform accruing to the individual may be diminished at the household level.

In this paper we analyse a UK educational reform, affecting all cohorts of individuals born after a specific threshold date of birth, which induced exogenous variation in the probability that a person obtained an academic qualification. By inducing a permanent shock to the qualifications distribution in affected cohorts, the reform induced a temporary shock to the cross-cohort qualifications composition for those individuals born in the neighbourhood of the threshold date. In absence of a preferred age gap there would be no gender-imbalance between potential partners in the marriage market. However, a prevailing age gap implies that women in the first cohorts affected by the reform are proportionately more qualified relative to their potential marriage partners born prior to the threshold. Therefore they will not be able to maintain both the typical spousal age gap as well as the same degree of positive assortative matching on education. This exogenous shock to the available choices of potential brides and grooms provides a unique opportunity to study the marriage market response to a qualifications shock.

Our results indicate that women who increased their qualification status in response to the reform also increased their probability of forming a marriage with a lower age gap, in particular by marrying a husband who would have also been affected by the reform. The observed decrease in the spousal age gap of 2.5 months is substantial compared to the sample mean age difference at marriage, and the results also indicate that affected women are not able to achieve the same degree of positive sorting on qualifications as before. We find a corresponding but weaker effect for men, who tend to slightly increase the age gap. Thus while treated women have to accept matches that deviate from the typical pattern in one or the other way, facing the choice of marrying younger or less qualified males, men obtain a better position at the marriage market as without changing the age gap men in the last unaffected cohort would find that their prospective partners were more qualified, and early treated male cohorts can increase assortative matching compared to men born just before the threshold, as their potential partners are from younger cohorts and thus also subject to the increased schooling requirement.

The marital age gap and its roots have attracted some attention in the literature. The economic explanation proposed by Bergstrom and Bagnoli (1993) postulates that males and females differ fundamentally in terms of economic prospects. The male breadwinner reveals his attractiveness in terms of income in later life, whereas the female's desirability is known from the start. Grooms with higher prospects have an incentive to postpone marriage until their high attractiveness is revealed, in order to be accepted by highly attractive brides. This difference in economic prospects produces the age gap. In contrast, Díaz-Giménez and Giolito (2013) advance a biological interpretation which ascribes the age gaps to differences in life-time fecundity profiles between genders. As female fecundity diminishes earlier than that of males, brides are inclined to accept marriage proposals at a young age, while grooms take the liberty of waiting.

Assortative matching on education due to complementarities in marital output (Becker, 1973, 1974) has become a well-established phenomenon in the marriage literature. However little is known regarding the relative strength of preferences over age and education in the formation of marital matches. Mansour and McKinnish (2013) argue that highly educated individuals meet similar-aged partners whilst in college and therefore ensuing marriages involve small age differences between spouses, whereas marriages with substantial age gaps are negatively selected as they are more likely to involve at least one spouse with low educational attainment. Holmlund (2008) analyses the impact on assortative mating and intergenerational mobility of a Swedish educational reform that not only increased the duration of compulsory education but also postponed ability tracking, with findings that indicate that although the reform led to sizeable increases in intergenerational income and educational mobility the degree of positive sorting between spouses increased, with assortative matching on education more important for women then men.

A substantial tranche of the demography literature has examined the effect of the 'marriage squeeze' on marital matches. The intuition is that sustained population growth implies an increase in the size of cohorts over time, and therefore with a preferred age gap between spouses there is a gender-imbalance in the marriage market due to an excess supply of age-appropriate women in comparison to men. Bronson and Mazzocco (2012) present evidence that birth cohort size is positively related with marriage rates but negatively associated with the age difference between spouses. Bhaskar (2012) in a theoretical model, shows that the spousal age gap does not respond to persistent population growth, but is the margin of adjustment to accommodate marriage market gender imbalances induced by transitory shocks to cohort size. Empirically, Bergstrom and Lam (1989) find that changes in gender ratios due to substantial fluctuations in fertility rates in Sweden in the early 20th century were largely accommodated by movement in the age difference between spouses. The 1958-1961 famine in China reduced cohort sizes substantially, however Brandt, Siow, and Vogel (2009) find that marriage rates were largely unaffected due to adjustments in the marital age gap.

Our results add a new aspect to the large literature on the causal effects of education that make use of compulsory school-leaving age reforms for identification. To our knowledge this is the first paper to investigate the effect of such a reform on matching behaviour in the marriage market along both the education and spousal age difference dimensions. A number of analyses have used the UK educational reforms to estimate the impact of education on various adult outcomes such as earnings (see, *inter alia*, Harmon and Walker (1995), Oreopoulos (2006) and Devereux and Hart (2010)), adult health and mortality (Clark and Royer, 2013), life-time wealth and happiness (Oreopoulos, 2007) and crime (Machin et al., 2011), as well as the intergenerational causal effect of parental education on outcomes of their children, such as health (Doyle, Harmon, and Walker, 2005; Lindeboom, Llena-Nozal, and van der Klaauw, 2009), or education (Galindo-Rueda, 2003; Chevalier, 2004; Chevalier et al., 2004). Our findings suggest a reform that is cohort-based and formally gender-neutral may have asymmetric impacts by gender via the marriage market. This may in turn have important implications for long-term outcomes that are more heavily dependent on the household environment rather than on the individual characteristics alone.

The remainder of the paper is structured as follows: Section 4.2 outlines the institutional context, the data used in the analysis and our empirical methodology are described in Section 4.3. We present and discuss our results in Section 4.4. Section 4.5 concludes.

4.2 Institutional Context

Compulsory schooling has been a feature of the education system in England and Wales since the late nineteenth century. Children are required to start education no later than the beginning of the academic year (September 1st - August 31st) after which they turn 5, and are required to remain in education until they have reached the legislated minimum school-leaving age. There are two tiers[2] of school-age academic qualifications: the first level of examinations is taken at the end of the academic year in which an individual turns

[2]Ordinary Levels (O'Levels), targeted towards academically inclined students and a pre-requisite for participation in further education, were introduced in the 1950s as the main academic qualification achieved at school. After the introduction of comprehensive schools the Certificate of Secondary Education (CSE) qualification was introduced in 1965 to meet the needs of the less academically-able. Both these exams are taken at age 16 and constitute the first tier of school qualifications. O'Levels and CSEs were replaced by a single examination, the General Certificate of Secondary Education (GCSE), in 1988. The second tier of qualifications are Advanced Levels (A'Levels), sat at age 18.

16; for more academically able students a second tier of qualifications are sat after two years of further study.

4.2.1 The reform

Since the introduction of the first minimum school-leaving age legislation in 1880 there have been a number of increases to the age until which students are compelled to remain in full-time education[3]. We focus on one of these increases, the Raising of School Leaving Age (RoSLA), which raised the schooling requirement by one year, from age 15 to age 16. The intention to implement RoSLA was first announced by the UK Government in 1964 and enacted in September 1972, affecting the mandatory school-leaving age of all individuals born after September 1st 1957.

Figure 4.1: RoSLA effect

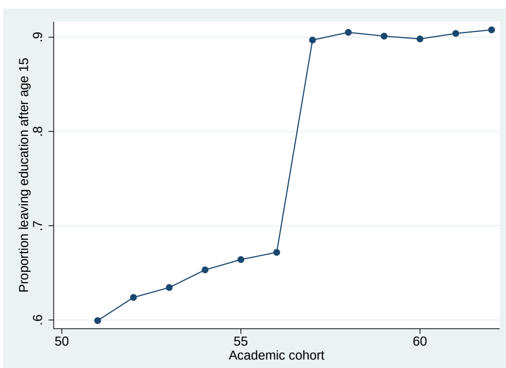

(a) Participation

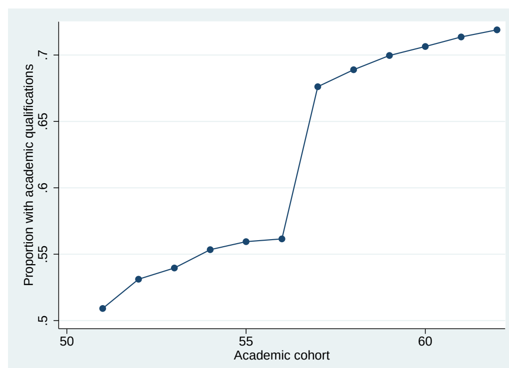

(b) Academic Qualifications

Notes: The graphs show the proportion of individuals a) participating in education after age 15 and b) with an academic qualification by academic cohort of birth (Sept-Aug). The RoSLA affected individuals born after after September 1st 1957.

The RoSLA reform impacted the school-leaving age of a substantial fraction of the population. Figure 4.1(a) shows that the proportion of individuals leaving education after age 15 increased approximately 25 percentage points in response to the new leaving age requirement. Furthermore, in comparison to other legislative increases to the minimum schooling requirement, RoSLA has the unique feature insofar that it raised compulsory schooling precisely up to the age at which the first tier of academic examinations are

[3]The first compulsory leaving age of 10 years was introduced by the Education Act (1880), raised to age 11 by the Elementary (School Attendance) Act (1893), increased again to age 14 in 1918 by the Fisher Act. The Butler Act (1944), initially raised the minimum leaving age to 15, but made provision for a subsequent rise in 1972 up to age 16. More recently the Education Act (2008) has introduced the Raising of Participation Age (RPA), which from September 2015 requires all individuals in England and Wales to remain in formal education or training until their 18th birthday.

taken. Thus by compelling students to stay in school for an additional year, the RoSLA reform induced an increase in the likelihood of them taking the examinations, and thereby increasing the probability of achieving a qualification. Figure 4.1(b) indicates that the introduction of the RoSLA increased the proportion of individuals obtaining an academic qualification increased by just over 10 percentage points. Chevalier et al. (2004) show that RoSLA's impact on qualifications was limited to the first tier of academic qualifications only, with no ripple-upward effects observed on higher qualifications.

To examine the robustness of our assertion that the marriage market responded to a temporary gender-age-qualifications imbalance induced by the RoSLA we examine another institutional rule in the English education system which induced exogenous variation in the propensity to receive a qualification *within* a cohort rather than *across* cohorts.

4.2.2 The Easter Leaving Rule

The Education Act (1962) introduced the Easter Leaving Rule (ELR) in response to concerns that if school-leaving eligibility is determined by the precise birth date alone, individuals born at the beginning of the academic year may not complete as much secondary education as later-born individuals, and therefore may be disadvantaged in the labour market due to the lower investment in human capital. The ELR imposed that persons born between September 1st and January 31st should remain in education until the end of the Spring term of the academic year in which they reached the compulsory school-leaving age, whereas those born between February 1st and August 31st were required to stay in school until the end of the Summer term.

Figure 4.2: ELR effect

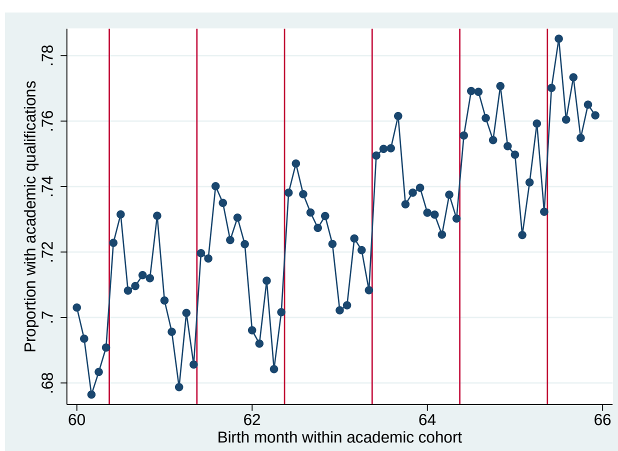

Notes: The graph displays the proportion of individuals holding an academic qualification by month of birth within an academic cohort. The vertical lines indicate the threshold of the Easter Leaving Rule (February within each academic cohort).

For post-RoSLA cohorts, who are required to remain in school until 16, which is the precise age at which the first-tier academic examinations are available, this implies that individuals within the same cohort have different probabilities of obtaining a qualification, as those born towards the beginning of the academic cohort could leave school before the examination period, whereas persons born in the latter part of the year, through the requirement to stay in school until the end of the Summer term, would have a higher propensity of sitting the examination, and consequently obtaining a qualification. Figure 4.2 shows that individuals born immediately after the ELR threshold have up to a 5 percentage point higher propensity of holding a qualification in comparison to individuals eligible to leave school prior to the examination period.

4.2.3 Why does the RoSLA induce a temporary imbalance in the marriage market?

We argue that by being applied at the cohort level the reform induced a temporary shock to the cross-cohort composition of qualifications by gender in the neighbourhood of the RoSLA threshold date. With a prevailing age gap between spouses of more than one year this implies that individuals form matches across academic cohorts. Specifically a typical woman will match with a man from an older cohort. As the RoSLA introduced a single threshold date applicable to all individuals regardless of gender, and matches tend to occur across cohorts, this implies that each side of potential matches in the neighbourhood of the threshold are differentially affected as there is a higher proportion of qualified individuals in post-RoSLA compared to pre-RoSLA cohorts. To clarify the imbalance we assess the potential age-qualifications matches in the vicinity of the threshold for each gender:

Women born in the RoSLA cohort:

In absence of the reform, women born in academic cohort 1957 or later would typically match with men born in 1956 or earlier. By increasing the fraction of individuals holding a qualification, the RoSLA increases the ratio of qualified women to qualified men thereby creating a gender imbalance in the qualifications composition across cohorts. A woman maintaining the typical age gap will face an increased likelihood of matching with a man who has lower qualifications than herself. In contrast a woman maintaining the typical sorting on qualifications will face an under-supply of appropriately qualified men in the usual cohort, therefore the attractiveness of younger men from post-RoSLA cohorts increases as they are proportionately more qualified than pre-RoSLA men such that they may become more acceptable as a potential match.

Men born in the RoSLA cohort:

In absence of the reform, men born in academic cohort 1957 or later would typically match with women born in 1958 or later. As both cohorts are subject to the increased education requirement of the RoSLA, there is no imbalance in the relative proportion of qualified men to qualified women.

Bhaskar (2012) argues that the marital age gap is the margin of adjustment by which the marriage market accommodates gender imbalances. The degree of adjustment will depend on the relative preferences of individuals over the age and qualifications of their partner. The larger the preference for qualifications, the greater the proportion of individuals who marry outside of the typical partner-age cohort resulting in a large adjustment in the age gap in the directly affected cohort. In contrast if age considerations are more important, or if the adjustment is spread out over several cohorts, then the direct effect on the age gap will be small, and the imbalance will be accommodated through qualification differences between spouses.

4.3 Empirical strategy

We empirically investigate how the marriage market responded to the temporary gender imbalance induced by the RoSLA and the prevailing age gap by examining outcomes around the threshold of the reform using a regression discontinuity design. We use a sample which encompasses the typical age gap either side of the discontinuity, therefore the method essentially summarises marriage market behaviour across cohorts.

4.3.1 Estimation Method

The RoSLA reform introduced a threshold date of birth according to which the minimum length of compulsory schooling was determined: individuals born prior to September 1957 were able to leave school at age 15, whereas those born after August 1957 were required to remain in education until 16 years of age. The reform can therefore be considered as a natural experiment, providing an exogenous source of variation to an individual's educational characteristics. As this variation was solely determined by an observable characteristic, the individual's time of birth, a regression discontinuity (RD) design is particularly suited as an estimation method in our analysis, where we explore responses to the RoSLA in the marriage market.

In essence the RD approach is based on the simple intuition that individuals born in the neighbourhood of 1st Sept 1957 are identical apart from which side of the threshold date they are born, and in absence of the RoSLA 'treatment' these individuals would have similar outcomes. Although the underlying distribution of the running variable is continuous, our dataset contains only discrete measures: month and year of birth. We therefore adopt the RD estimation procedure proposed by Lee and Card (2008). Visual depictions of the results are produced using local polynomial smooths, whereas the base specification for our analytical results is:

$$Y_{ij} = \alpha_0 + D_j\beta_0 + P_j^l\gamma_0 + (D_j \times P_j^l)\delta_0 + a_j + \epsilon_{ij} \tag{4.1}$$

where Y_{ij} is the outcome for individual i born in month j; P_j^l is a vector of polynomial functions in the running variable, x_j, with $(l = 1, 2, 3)$; D is an indicator of whether the individual was subject to RoSLA which is interacted with the polynomial to allow these to be different left and right of the discontinuity. We determine the appropriate window width of observations around the discontinuity to use in the estimation following the cross-validation procedure suggested by Ludwig and Miller (2007).[4]

4.3.2 Data

The Labour Force Survey (LFS) is the largest representative survey undertaken in the UK, with around 11,000 private households interviewed each wave. The survey contains detailed information on each individual within the household, such as their marital status, ethnicity, date of birth (measured by month and year), education, nationality and country of birth, as well as the relationships between each of the household members[5]. We pool data from the 1975-2006 surveys, the latter being the last data for which date of birth information is publicly available.

We define an individual's academic birth cohort by whether the individual was born between September 1st of year t and August 31st of year $t + 1$. Thus the first individuals subject to the RoSLA, aged 15 when they started the school year in September 1972, were born in academic cohort 1957. The spousal age gap is calculated as the linear difference between the husband's and wife's ages, measured in months, at the time of the survey. As the RoSLA did not induce an impact on educational achievement beyond the first tier of academic qualifications (Chevalier et al., 2004) we can, without loss of generality, focus on the binary outcome of whether an individual has an academic qualification or not. Respondents in the LFS are asked to record their qualifications according to their equivalence in a categorisation of English qualifications. From this information we are unable to ascertain whether individuals were educated in England and Wales. Therefore to mitigate the inclusion of individuals not subject to the relevant schooling system, we restrict the sample to those individuals born in the UK, but resident within England or Wales at the time of survey. As with the spousal age gap, we construct the marital

[4] A full description of the estimation techniques is presented in the technical annex.

[5] Unfortunately the LFS does not measure parity of marriage and has no information regarding the age at marriage. However this is only problematic for the analysis if spouses in higher parity marriages have different preferences over age and qualifications of their partners than spouses in first marriages. By imposing age restrictions to the sample we can mitigate the inclusion of higher parity marriages. This robustness check is explored in Appendix C.

qualification gap as the difference between the binary indicators of whether each spouse has a qualification. This variable is therefore equal to 1 if the husband has a qualification and the wife does not, is equal to 0 if both spouses have the same qualification status, and -1 if the wife has a qualification and the husband does not.

Table 4.1: Descriptive statistics

	N	Mean	S.D.	Min	Max
		Married Women			
Qualifications	116,709	0.68	0.47	0	1
White	138,133	0.99	0.09	0	1
Age	143,108	33.74	7.57	20	50
Survey year	143,108	91.11	7.67	75	106
Age gap	143,108	24.89	37.28	-120	120
Diff. qualif.	115,089	-0.08	0.53	-1	1
		Married Men			
Qualifications	108,965	0.65	0.48	0	1
White	126,020	0.99	0.09	0	1
Age	128,853	34.55	7.29	20	50
Survey year	128,853	91.86	7.45	75	106
Age gap	128,853	18.66	37.02	-120	120
Diff. qualif.	107,966	-0.07	0.53	-1	1

Notes: The table shows descriptive statistics of key variables. Sample restrictions: age gap is within 120 months to either side; own age is between 20 and 50 years; individuals born within academic cohorts 1951-1962.

We create a sample of men and a sample of women in order to analyse responses around the RoSLA threshold for each individual spouse, descriptive statistics are displayed in Table 4.1. In each sample we retain only prime-age individuals (aged 20-50) within academic cohorts of birth close to the RoSLA implementation (cohorts 1951-1962) and exclude the small number of couples where the spousal age difference is above or below 10 years. After restrictions, for the age gap analysis we have 128,853 (143,108) dyads where the man (woman) is born in relevant period. As detailed information regarding the level of qualifications held by an individual has been measured in the LFS only since 1984 the male and female samples used in the qualification gap analysis is reduced to 108,965 and 116,709 couples respectively.

4.4 Results

We maintain that in the presence of preferred age gaps between partners, the RoSLA reform induced a temporary imbalance in the proportion of qualified individuals across cohorts, and show evidence that this imbalance significantly impacted the composition of marital matches in the RoSLA-neighbourhood cohorts. We expect that qualified women from the first RoSLA cohorts would more frequently marry unqualified men, or reduce the marital age gap in comparison to their counterparts born prior to the RoSLA threshold as their choices are constrained as a result of the imbalance, which implies they are unable to maintain both the typical age gap and qualifications sorting. In contrast, we hypothesise that RoSLA treated men marry similarly educated partners compared to pre-RoSLA men. We examine the robustness of our assertion in two ways. First we examine another institutional rule in the English education system which induced exogenous variation in the propensity to receive a qualification within rather than across a cohort, second we examine between cohort thresholds in non-RoSLA years to confirm that our observed results are unique to the RoSLA reform discontinuity.

We first verify the effect of the RoSLA reform on an individual's qualification status, before examining our main results of how the spousal age and qualification gaps respond at the threshold. We then examine changes in match-types around the discontinuity to assess the extent of substitution between partner age and qualifications.

4.4.1 Qualifications

Figure 4.3: RD on own academic qualification by month of birth

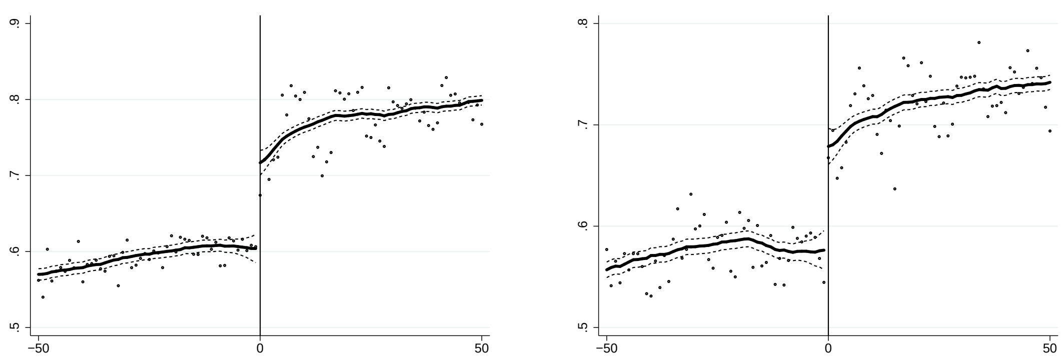

(a) Married Women (b) Married Men

Notes: The graphs show local polynomial smooths indicating the likelihood an individual has obtained an academic qualification. The dots reflect means by each bin on the abscissa. The solid line is the local polynomial smooth with the bandwidth and degree as shown using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

Figure 4.3 shows the proportion of (a) married women and (b) married men who hold academic qualifications by each respective spouse's distance of birth, in months, to the RoSLA threshold date which has been normalised to zero. The reform was associated with a substantial increase in the likelihood of obtaining a qualification, the impact for married women, 11.4 percentage points, slightly larger than that for married men, 10.2 percentage points[6], which are comparable to the findings of Dickson and Smith (2011) who find an impact of 9.5 percentage points in their estimation for working-age men.

4.4.2 Marital Age Gap

Figure 4.4 displays the marital age gap by each spouse's distance of birth, in months, to the RoSLA threshold date which has been normalised to zero. The estimates are produced using the optimal bandwidth of 24 months, as indicated by the Ludwig and Miller (2007) cross-validation procedure, and a quadratic polynomial of the running variable. The figure confirms the basic hypothesis: for married women at the RoSLA threshold we see a clear reduction in the age gap of approximately 2 months. In contrast for men there is a small, but insignificant, increase in the age gap.

Figure 4.4: Marital age gap at RoSLA threshold

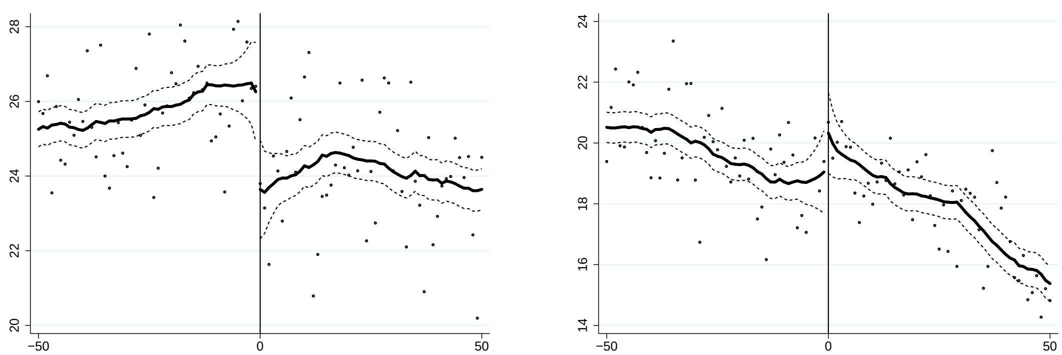

(a) Married Women (b) Married Men

Notes: The graphs show local polynomial smooths of marital age gaps in months (male-female). The dots reflect mean age gaps by each bin on the abscissa. The solid line is the local polynomial smooth with bandwidth 24 and degree 2 using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

Table 4.2 displays the analytical results associated with Figure 4.4. For each gender we first present the results of the estimation using our preferred bandwidth of 24 months, and explore the robustness of these estimates to using half and double our preferred bandwidth, and also for a linear, quadratic and cubic polynomial in the running variable over columns 1-3 respectively. In columns 4-6 we add a set of basic covariates (age and ethnicity) to the base specifications. We report the Lee and Card (2008) G-statistic along with the Akaike (AIC) information criterion to test the goodness of fit of the polynomial used. The upper panel displays the estimates for women revealing that for a given bandwidth the G-statistic does not clearly suggest a polynomial degree, but tends to favour a more

[6] The analytical results associated with Figure 4.3 are presented in Appendix C.1, Table C.1.

complex polynomial. Conversely, the AIC generally indicates that the smallest (linear) polynomial is appropriate. A second degree polynomial seems sensible as a compromise of these different test results, noting that Gelman and Imbens (2014) find that RD estimators for causal effects using higher-order polynomials can be misleading and suggest that, in absence of conclusive evidence in favor of using a complex polynomial, a linear or quadratic polynomial in the running variable yields more credible results. The G-statistic is even less suggestive in the lower panel for married men, which clearly indicates that a linear polynomial is preferred according to the AIC, however we note that there is no significant difference in the coefficients obtained using either a first or second degree polynomial with the optimal bandwidth.

The RD estimate for married women in our preferred specification (column (2) of the upper panel), which uses the optimal bandwidth of 24 months and a 2nd-order polynomial, indicates that the RoSLA induced a negative and statistically significant response on the age gap of 2.5 months, 10% relative to the sample mean. With other bandwidth choices the effect is somewhat amplified, a reduction of 3.9 (3.0) months for half (double) the optimal bandwidth, both statistically significant[7]. Including basic controls does not alter the results significantly, indicating there are no discontinuities in the covariates around the threshold. In the lower panel, the estimates for married men using a 2nd degree polynomial reveal a consistent but statistically insignificant response to RoSLA of 1.1 months over all bandwidth choices. With a linear polynomial the magnitude of the estimates increase somewhat and gain significance for non-optimal bandwidths, but remain significantly lower in magnitude than the response for women.

As the LFS does not record the parity of marriage, a potential concern is that the results would be biased if spouses in higher parity marriages have different sorting patterns over qualifications and age than spouses in first marriages. We explore robustness to this in two ways. First we mitigate the inclusion of higher parity marriages by applying upper age bounds (40, 45 and 50) to individuals included in the LFS sample, see Table C.3 in Appendix C.1, finding no significant differences in the estimates for males over the different samples, and qualitatively similar results between the age-restricted samples and the main analysis for female, with the magnitude of the response increasing as the upper age bound decreases. Second, we perform a parallel analysis on a 1% sample from the census, where we constrain the sample to include first marriages only, yielding results consistent with those obtained from our LFS samples[8], see Figure C.2 in Appendix C.1.

[7]We explore the sensitivity of the estimates to the choice of bandwidth further by performing the analysis for all bandwidths between 12 and 60 months, with results suggesting that the estimates are qualitatively robust and stable over a wide range of bandwidth choice. These results are available in the Technical Annex.

[8]The Census Longitudinal Study (CLS) is comprised of linked census records from the 1971-2001 censuses of individuals born on four specific days of the year, capturing approximately 1% of the population.

Table 4.2: Change in marital age gap at RoSLA threshold

	(1)	(2)	(3)	(4)	(5)	(6)
			Married Women			
Optimal Bandwidth (24)	-3.054***	-2.537**	-5.836***	-2.882***	-2.598**	-5.757***
	(0.695)	(0.825)	(1.326)	(0.683)	(0.797)	(1.300)
N	48,657	48,657	48,657	48,104	48,104	48,104
G-statistic (p-value)	0.022	0.016	0.089	0.018	0.011	0.059
AIC	491,926	491,929	491,922	486,128	486,132	486,125
1/2 x Optimal B. (12)	-4.595***	-3.898**	-0.674	-4.544***	-3.859**	-0.692
	(0.790)	(1.057)	(1.264)	(0.772)	(1.027)	(1.191)
N	24,273	24,273	24,273	24,116	24,116	24,116
G-statistic (p-value)	0.493	0.451	0.619	0.374	0.325	0.463
AIC	245,358	245,361	245,360	243,639	243,642	243,641
2 x Optimal B. (48)	-2.410***	-2.989***	-3.238***	-2.305***	-2.977***	-3.056***
	(0.548)	(0.711)	(0.775)	(0.546)	(0.703)	(0.754)
N	96,657	96,657	96,657	94,320	94,320	94,320
G-statistic (p-value)	0.013	0.012	0.011	0.011	0.010	0.010
AIC	975,294	975,297	975,299	951,731	951,733	951,736
			Married Men			
Optimal Bandwidth (24)	1.170	1.160	2.040*	1.216*	1.066	2.036*
	(0.594)	(0.697)	(0.855)	(0.597)	(0.710)	(0.846)
N	44,466	44,466	44,466	44,239	44,239	44,239
G-statistic (p-value)	0.823	0.871	0.850	0.734	0.804	0.782
AIC	446,900	446,901	446,904	444,292	444,292	444,295
1/2 x Optimal B. (12)	2.087**	1.066	-1.302	2.019**	1.019	-0.978
	(0.691)	(0.817)	(1.158)	(0.695)	(0.812)	(1.125)
N	22,373	22,373	22,373	22,314	22,314	22,314
G-statistic (p-value)	0.814	0.826	0.872	0.695	0.732	0.744
AIC	224,543	224,545	224,547	223,791	223,793	223,795
2 x Optimal B. (48)	1.668***	1.083	1.485*	1.704***	1.145	1.416*
	(0.443)	(0.646)	(0.712)	(0.445)	(0.645)	(0.712)
N	87,187	87,187	87,187	86,106	86,106	86,106
G-statistic (p-value)	0.456	0.430	0.449	0.462	0.433	0.451
AIC	876,672	876,675	876,677	865,532	865,535	865,536
Polyn. degree	1	2	3	1	2	3
Basic controls	No	No	No	Yes	Yes	Yes

Notes: The table shows estimates from local parametric estimation of equation (4.1) as described in Section 4.3.1 using different bandwidths, over rows, and polynomial degrees 1 to 3 over columns. The dependent variable is the spousal age gap measured in months (male-female). The bandwidth reflects the number of values of the running variable (month of birth) on each side of the discontinuity. Standard errors are robust and allow for random and identical specification errors. Below the estimates the p-value of the Lee and Card (2008) G-statistic and the Akaike Information Criterion indicate goodness of the polynomial fit. Robust standard errors reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

To preserve confidentiality the data released for analysis is restricted to academic year of birth for each spouse only. Although the CLS sample is larger (214,325 dyads) than our LFS dataset, we retain the LFS in our main analysis as the RDD framework is more credible with a fine discretisation of the running variable (Lee and Lemieux, 2010).

4.4.3 Marital Qualification Gap

Figure 4.5 displays the spousal qualifications gap around the RoSLA discontinuity[9]. As described in Section 4.3.2 a qualifications gap equal to 1 involves a qualified husband and unqualified wife, a gap of 0 indicates equally qualified spouses and a gap of -1 results when the wife is qualified but her husband is not. For married women the response occurs contemporaneous to the reform implementation, women affected by RoSLA clearly increase their qualification relative to their husbands, reflected in the negative shift of the spousal qualification gap. The magnitude of the difference declines over subsequent cohorts before leveling off. We see that after leveling off the post-RoSLA qualifications difference is greater than the pre-RoSLA difference, consistent with the increase in qualifications at the RoSLA being larger for women as indicated by Figure 4.3. In contrast, the response in the qualifications difference for married men happens in the immediate cohorts prior to RoSLA, as their younger spouses are increasingly likely to be born after the threshold and are therefore are proportionately more qualified. For men born after the RoSLA threshold the qualifications difference is restored as their younger potential partners are also subject to the increased schooling requirement ending the marriage market imbalance.

Figure 4.5: Marital qualifications gap at RoSLA threshold

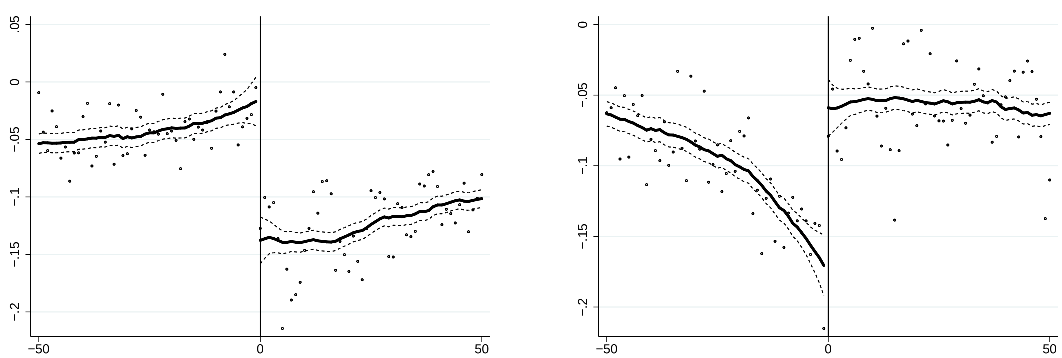

(a) Married Women (b) Married Men

Notes: The graphs show local polynomial smooths of difference in qualifications between spouses (male-female). The dots reflect mean qualifications difference by each bin on the abscissa. The solid line is the local polynomial smooth with bandwidth 24 and degree 2 using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

So far the results have shown that the marriage market adjusts along both the age gap and the qualifications gap. For women passing the RoSLA threshold their chances of matching with an equally qualified husband from an older cohort decreases. Correspondingly, men born before the threshold find fewer equally qualified women in younger

[9]The analytical results associated with Figure 4.5 are presented in Appendix C.1, Table C.2.

cohorts and passing the threshold relieves this constraint. In order to examine the degree of substitution between the two dimensions we examine the proportion of match types at the discontinuity.

4.4.4 Substitution between Age and Qualification

In order to reconcile our findings with regard to the responses of the marital age and qualifications gaps we examine the degree of substitution between the two. We summarise the marriage responses around the RoSLA threshold into binary variables indicating the relative qualification level between spouses and whether an individual's spouse is from the pre- or post-RoSLA regime.

Figure 4.6: Spouse characteristics round Rosla

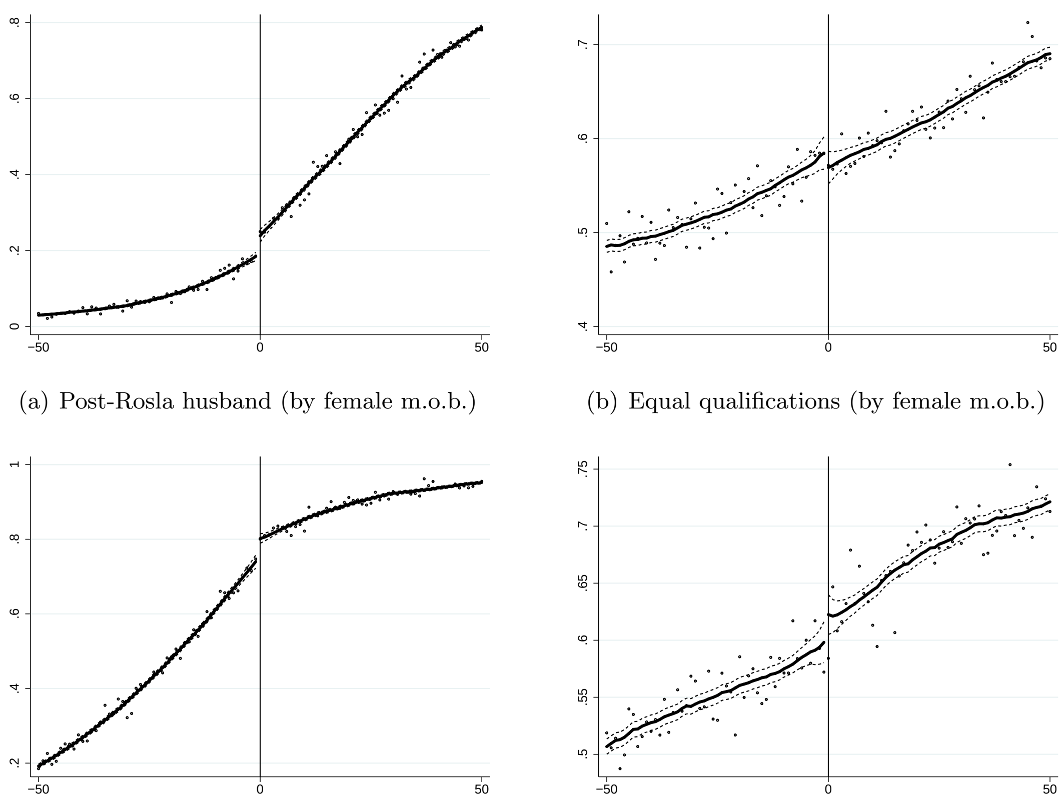

(a) Post-Rosla husband (by female m.o.b.)

(b) Equal qualifications (by female m.o.b.)

(c) Post-Rosla wife (by male m.o.b.)

(d) Equal qualifications (by male m.o.b.)

Notes: Graphs (a) and (c) show local polynomial smooths of whether a spouse is from a pre- or post-RoSLA cohort; graphs (b) and (d) display the proportion of marriages involving equally qualified partners. The dots reflect mean qualifications difference by each bin on the abscissa. The solid line is the local polynomial smooth with bandwidth 24 and degree 2 using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

Figure 4.6 considers the characteristics of partners around the discontinuity with regard to qualifications and age separately. For both genders at the threshold there is an observable jump in the proportion of spouses from the post-RoSLA regime. For females there is a slight decreases in assortative matching, conversely for males the proportion of equally qualified spouses increases somewhat.

To examine the degree of substitution we combine the binary indicator of whether an individual has a spouse from a pre- or post-RoSLA cohort with the indicators of the relative qualification. The combination of these indicators yields a choice set of six options for each spouse at the discontinuity: 1) a match with a pre-RoSLA spouse where the husband is qualified but the wife is not, 2) a pre-RoSLA spouse where both spouses are equally qualified, 3) a pre-RoSLA spouse where the wife only has an academic qualification, 4) a post-RoSLA spouse where only the husband is qualified, 5) a post-RoSLA spouse where the husband is not qualified and the wife has an academic qualification and 6) a post-RoSLA spouse who is equally qualified.

Figure 4.7: Match types of married women

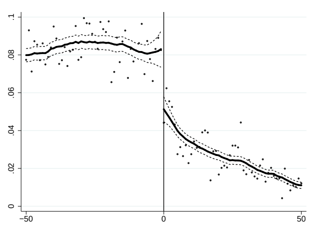

(a) Pre-RoSLA, more qualified

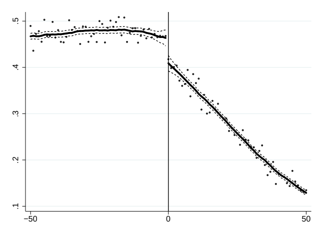

(b) Pre-RoSLA, equally qualified

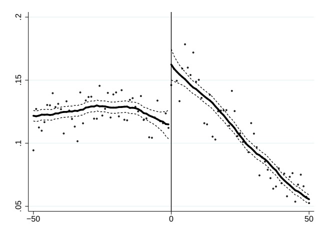

(c) Pre-RoSLA, less qualified

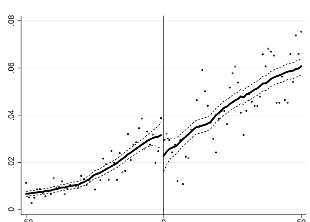

(d) Post-RoSLA, more qualified

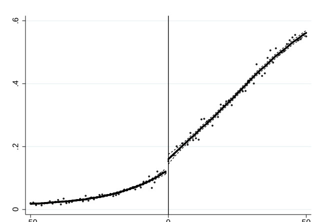

(e) Post-RoSLA, equally qualified

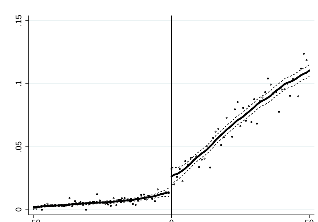

(f) Post-RoSLA, less qualified

Notes: The graphs show local polynomial smooths of match types with regard to spousal characteristics. The dots reflect mean age gaps by each bin on the abscissa. The solid line is the local polynomial smooth with bandwidth 24 and degree 2 using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

Figures 4.7 and 4.8 shows the marriage market matches respectively of women and men around the RoSLA threshold. For both genders the predominant matchings occur between spouses with equal qualifications ((b) and (e)), women with older (pre-RoSLA men) and men with younger (post-RoSLA women).

As shown in Section 4.4.1 at the discontinuity the proportion of women obtaining an academic qualification increases, therefore unsurprisingly we observe that the match between a wife and a pre-RoSLA husband who is more or equally qualified clearly drops at the discontinuity (Figure 4.7(a) and (b)), however the corresponding increase in matches between wives and less qualified pre-RoSLA husbands is not large enough to fully compensate the decrease, indicating that adjustment on the qualification dimension alone does not adequately explain the observed sorting patterns. We see small but significant increases in the proportion of women matching with younger (post-RoSLA) men with equal or less qualifications than themselves, which outweigh the smaller decrease in the proportion of matches involving a more qualified post-RoSLA husband.

Figure 4.8: Match types of married men

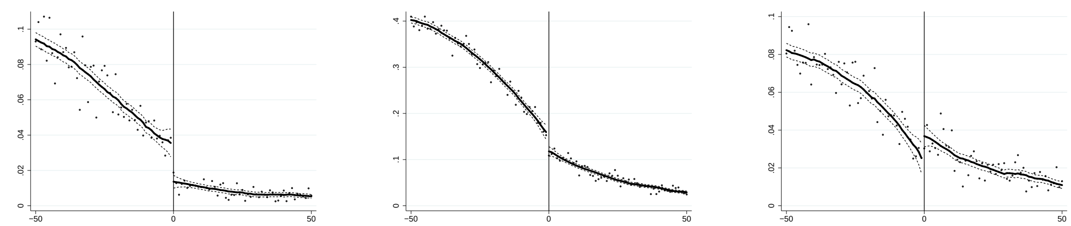

(a) Pre-RoSLA, more qualified (b) Pre-RoSLA, equally qualified (c) Pre-RoSLA, less qualified

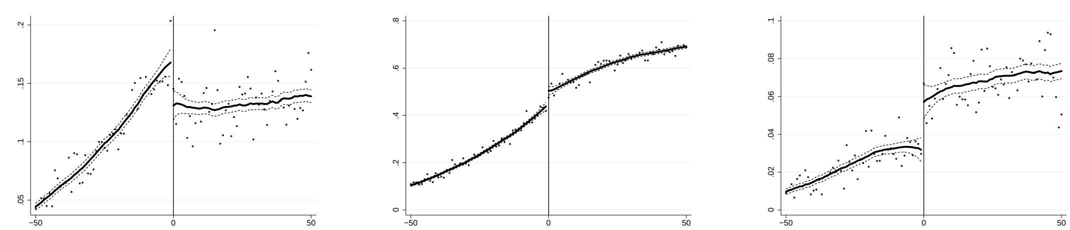

(d) Post-RoSLA, more qualified (e) Post-RoSLA, equally qualified (f) Post-RoSLA, less qualified

Notes: The graphs show local polynomial smooths of match types with regard to spousal characteristics. The dots reflect mean age gaps by each bin on the abscissa. The solid line is the local polynomial smooth with bandwidth 24 and degree 2 using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

The response for men at the discontinuity, Figure 4.8, is similar to that observed for women. The most pervasive match type is with a post-RoSLA equally qualified spouse (Figure 4.8(e)), which clearly increases at the threshold. As the proportion of men with an academic qualification rises at the threshold, this increase is driven by an increase in matches where both spouses hold an academic qualification. There is also a marked increase in the proportion of matches involving husbands who are more qualified than their wives, which is consistent with Fisman, Iyengar, Kamenica, and Simonson (2006), who find in a speed-dating experiment that men prefer women with a level intelligence that does not exceed their own.

Table 4.3: Change in match types around the discontinuity

	Women			Men		
Husband's qualification relative to wife's	More	Equal	Less	More	Equal	Less
Pre-RoSLA Spouse	-0.033***	-0.053***	0.048***	0.014**	-0.033***	-0.021***
	(0.007)	(0.008)	(0.009)	(0.005)	(0.008)	(0.004)
N	48657	48657	48657	44466	44466	44466
Post-RoSLA Spouse	-0.009	0.033***	0.012**	0.026***	0.055**	-0.039**
	(0.006)	(0.008)	(0.004)	(0.006)	(0.018)	(0.018)
N	48657	48657	48657	44466	44466	44466

Notes: The table shows estimates from local parametric estimation of equation (4.1) as described in Section 4.3.1 using the preferred specification (bandwidth of 24 months and a quadratic polynomial in the running variable. The dependent variable is the relative qualification level of the spouse, over columns, estimated separately for pre-and post-RoSLA cohorts, over rows. Standard errors are robust and allow for random and identical specification errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The analytical results associated with Figures 4.7 and 4.8 are displayed in Table 4.3. The upper left quadrant characterises the matches of women with older spouses, who are not subject to the RoSLA and in absence of the reform would constitute the typical partner for a woman born in academic year 1957. At the discontinuity the proportion of women holding a qualification increases relative to pre-RoSLA men, and therefore the proportion of matches where the husband is more or equally qualified decreases mechanically. The estimates suggest that 56% (0.048/(0.033+0.053)) choose to retain the positive age gap with a less qualified husband, whereas 38% (0.033/(0.033+0.053)) sacrifice the age gap in order to maintain homogamy in education. We also observe that 14% (0.012/(0.033+0.053)) of women at the threshold retain neither the preferred age nor qualifications match, which may reflect general equilibrium effects or other unobserved factors which gain importance if a husband of the preferred age-qualification type is not available. The corresponding estimates for men are presented in the right hand panel of Table 4.3, where the lower quadrant would in absence of the reform represent the typical matches for early reform-treated men. At the threshold the proportion of men with qualifications increases restoring the typical gender-qualifications composition and the proportion of matches where the husband is less qualified decreases mechanically. The degree of homogamy increases, as both potential partners are subject to the increased schooling requirement. Overall the proportion of matches with older women decreases, but for the share of marriages involving a more qualified husband.

This latter analysis indicates that women care more about the age of their spouse than his level of qualifications. For men the relative importance seems balanced as the increase in the preferred match by RoSLA-affected men originate in equal shares from substituting matches with older pre-RoSLA women and younger women with higher qualifications. This finding contrasts with Belot and Francesconi (2013) who examine partner choices in a speed-dating context, finding that women and men put comparable weights on physical attributes, such as age, and that education of potential partners has only minor importance in shaping mate selection, albeit given greater worth by males. However these dissimilarities may be driven by the differences in partner selection in the dating as opposed to the marriage market.

4.4.5 Robustness

To assess the validity of our argument, that the response is induced by an imbalance in age-qualification types across cohorts, we examine another feature of the UK educational system, the Easter Leaving Rule (ELR), which increased the probability of qualifications within an academic cohort, but not across cohorts. We then verify that our analysis is unique to the RoSLA threshold rather reflecting typical between cohort effects. We repeat the analysis with a placebo discontinuity, examining the threshold four years earlier (1952/1953) where there is no cross-cohort gender imbalance in qualification as both pre- and post-threshold cohorts were subject to the pre-RoSLA schooling regime. Furthermore, we apply a difference-in-difference approach in the context of the regression discontinuity design (RD-DID), using the placebo cohorts to form counterfactual observations.

4.4.5.1 A within-cohort increase in qualifications

Using the exogenous source of variation in qualifications induced by the ELR, which is independent of variation in the duration of schooling, Anderberg and Zhu (2014), analysing outcomes of married women, find that the ELR induced an unambiguous increase in the probability of obtaining an academic qualification; furthermore that women born after the ELR threshold are more likely to have qualified husbands relative to women born earlier in the academic year. In order to examine the qualification effect independent of the age effect for both spouses we examine the outcomes for individuals from post-RoSLA cohorts (1959-1966) only, such that with the prevailing age gap both partners would be subject to the RoSLA treatment and therefore there should be no imbalance in age-qualifications types. Figure 4.9 displays the marital age gap by each spouse's distance of birth, in months, to the ELR January-February threshold, normalised to 0. The graph confirms that the within cohort variation in qualifications induced by the ELR does not result in a discontinuity in the age gap of spouses around the January-February threshold, thereby supporting our assertion that qualifications do not exert an impact on the marital age gap in absence of cross-cohort variation in qualifications.

Figure 4.9: Placebo test of education reform w/o scarcity of types

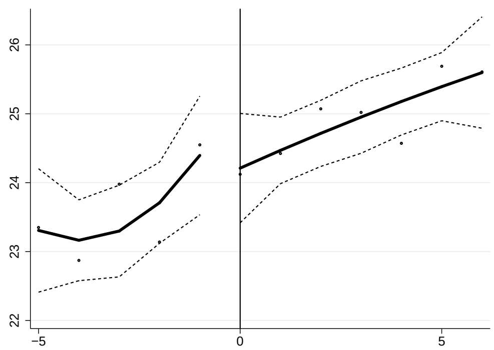

(a) Wives

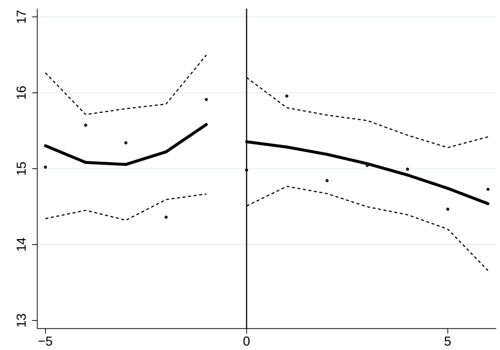

(b) Husbands

Notes: The graphs show local polynomial smooths of age gaps in months (husband-wife). The running variable is calendar month of birth. The dots reflect means by each bin on the abscissa. The solid line is the local polynomial smooth with bandwidth 24 and degree 1 using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

4.4.5.2 Robustness to inherent between-cohort differences

When using the regression discontinuity design approach the crucial identifying assumption is that individuals in the neighbourhood of the discontinuity are identical in characteristics. At the limit the comparison is between individuals born at the end of one academic cohort and the beginning of the next, and there may be an inherent degree of marital sorting across cohort thresholds that occurs in all years. To verify that our results are unique to the RoSLA threshold we repeat the analysis using a placebo year threshold, four years prior to the actual reform implementation. Figure 4.10 displays the marital age and qualifications gap by each spouse's distance of birth, in months, to the threshold (September 1953). As with the main results presented in Sections 4.4.2 and 4.4.3 the estimates are produced using the optimal bandwidth of 24 months and a quadratic polynomial of the running variable. The figures indicate no discontinuities in the qualifications difference between spouses, however there is evidence that the age gap adjusts at the threshold, although the difference at the threshold is not statistically significant.

In order to account for any inherent between cohort differences at the August-September threshold in non-reform years we use a difference-in-difference procedure similar to that proposed by Danzer and Lavy (2013), where the placebo cohorts above are used to form the counterfactual observations. We estimate the following specification:

$$Y_{ij} = \alpha_0 + \beta_0 After_{ij} + \beta_1 Rosla^{*} After_{ij} + \psi Period + \gamma_0 P_j^l + \delta_0 (T_i \times P_j^l) + a_j + \epsilon_{ij} \quad (4.2)$$

where Y_{ij} is the outcome of interest for individual i born at a distance of j from the relevant

Figure 4.10: Placebo Analysis

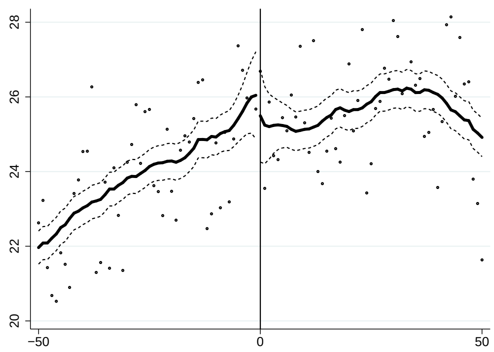

(a) Marital Age Gap (by female m.o.b.)

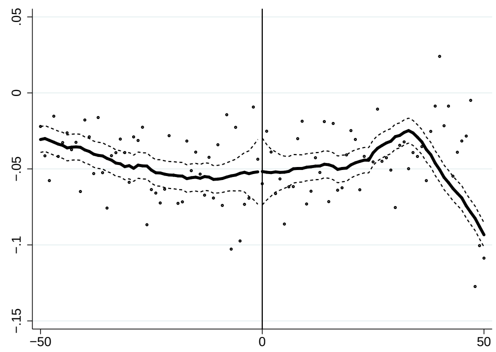

(b) Qualifications Gap (by female m.o.b.)

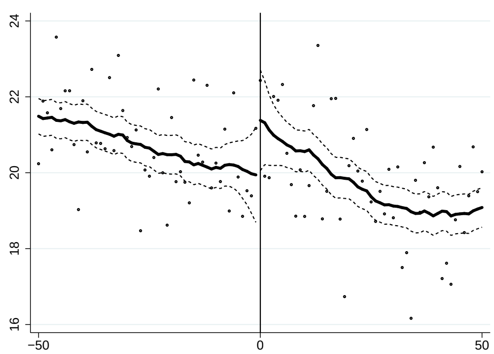

(c) Marital Age Gap (by male m.o.b.)

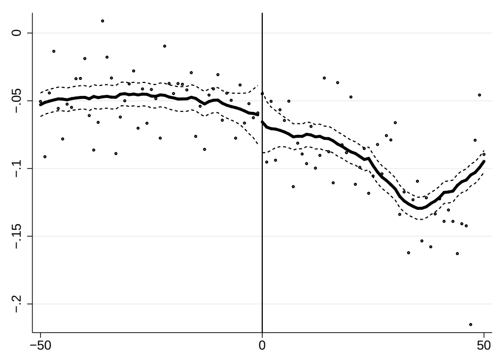

(d) Qualifications Gap (by male m.o.b.)

Notes: Graphs (a) and (b) show local polynomial smooths of the marital age and qualifications gap by female month of birth relative to the RoSLA threshold. Graphs (c) and (d) display the analogous plots for males. The dots reflect mean qualifications difference by each bin on the abscissa. The solid line is the local polynomial smooth with bandwidth 24 and degree 2 using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial.

threshold; $After$ is an indicator variable denoting the individual is born after the relevant discontinuity; $Rosla^{*}After$ is a dummy equal to 1 if the observation is after the threshold in the period around the RoSLA discontinuity. Therefore the RD-estimate is described by β_1. As in the base specification, equation (4.1), γ and δ capture the polynomial smooth in the running variable, a period dummy is introduced for each window of observations.

Table 4.4 compares the results using the regression discontinuity around the RoSLA those produced in the difference-in-difference approach. Considering the marital age gap, the comparison reveals that there are no significant disparities between the estimates produced by either procedure, whereas with the qualifications gap the effect of the RoSLA reform is slightly muted but remains significantly different from zero. Table 4.5 applies the RD-DiD to examine the substitution between marital age and qualifications gaps around

Table 4.4: Difference-in-Difference: Age and Qualifications

	Women		Men	
	Age Gap	Quals Gap	Age Gap	Quals Gap
RoSLA RD	-2.534**	-0.123***	1.159	0.117***
	(0.825)	(0.019)	(0.698)	(0.025)
N	48,668	39,775	44,475	38,050
RD-DiD	-2.881***	-0.100***	0.792	0.084***
	(0.805)	(0.015)	(0.504)	(0.018)
N	104,966	80,099	95,941	76,936

Notes: The upper panel shows estimates from local parametric estimation of equation (4.1) as described in Section 4.3.1 around the RoSLA threshold, using the preferred specification (bandwidth of 24 months and a quadratic polynomial in the running variable). The lower panel displays estimates from the regression discontinuity difference-in-difference procedure as described by equation 4.2 in Section 4.4.5.2. Standard errors are robust and allow for random and identical specification errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

the RoSLA threshold. The results are qualitatively similar to those presented in Table 4.3. Overall, the evidence presented in Tables 4.4 and 4.3 indicate that our main results remain robust after controlling for possible underlying between-cohort effects.

Table 4.5: Change in match types around the discontinuity

	Women			Men		
Husband's qualification relative to wife's	More	Equal	Less	More	Equal	Less
Pre-RoSLA Spouse	-0.023***	-0.020**	0.048***	0.004	-0.034***	-0.020***
	(0.005)	(0.009)	(0.009)	(0.004)	(0.006)	(0.003)
N	104,966	104,966	104,966	95,941	95,941	95,941
Post-RoSLA Spouse	-0.007	0.027***	0.006	0.012**	0.071***	-0.033**
	(0.005)	(0.006)	(0.004)	(0.006)	(0.018)	(0.012)
N	104,966	104,966	104,966	95,941	95,941	95,941

Notes: The table shows estimates from the regression discontinuity difference-in-difference specification, equation 4.2 in Section 4.4.5.2, using the preferred specification (bandwidth of 24 months and a quadratic polynomial in the running variable. The dependent variable is the relative qualification level of the spouse, over columns, estimated separately for pre-and post-RoSLA cohorts, over rows. Standard errors are robust and allow for random and identical specification errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

4.5 Conclusion

In this paper we have examined two pervasive features of the majority of marriages: husbands are on average older than their wives, and spouses tend to have the same level of education. We have investigated the impact of an educational reform, the Raising of School Leaving Age (RoSLA) implemented in England and Wales in September 1972, which induced a transitory imbalance in the cross-cohort qualifications composition, in order to elicit the relative strengths of preferences between age and qualifications.

We find that the reform induced a decrease in the spousal age gap of women who were induced into obtaining a qualification, the estimated reduction of 2.5 months is substantial compared to an average age gap of 24.8 months. Furthermore, affected women are not able to achieve the same degree of positive sorting on qualifications and are forced to accept atypical matches. In particular, affected women often choose lower qualified husbands with a preferred age gap and in fewer cases younger husbands from treated cohorts. We find a corresponding but reversed effect for affected men, who are able to return to the prevailing sorting patterns as their potential partners in the ideal age are from younger cohorts and would therefore also be subject to the increased schooling requirement. In contrast, men born just before the threshold face the transitory imbalance across cohorts. As with women born after the RoSLA threshold, these men cannot achieve the typical match characteristics as women in their ideal age range are already treated. In consequence, men born just before the threshold who have to deviate from the typical match choose in equal shares women from older cohorts or more qualified women with the preferred age gap.

Our results have potential implications for analyses that use changes to compulsory school leaving age legislation to elicit causal effects of education. Any cohort specific reform that is beneficial individually receives gender heterogeneity via the marriage market. Treated women are disadvantaged via the marriage market because of the positive age gap and their tendency to marry men from not yet affected cohorts. In an RD setting these women are compared to those just before the threshold who can have the typical match. Treated men, on the other hand, can achieve the typical match on the marriage market as their potential younger partners are already treated and they benefit from the reform at the same time. Compared to men born just before the threshold they are advantaged both by the reform and via the marriage market. Therefore, any observed gender heterogeneity in reform effects may be driven through the marriage channel rather than by the change in education. Moreover, this is especially important for long-term outcomes that are more heavily dependent on the household environment rather than on the individual characteristics alone.

One important caveat to our analysis is that by estimating the effects through a regression discontinuity design, our estimates are applicable only to those individuals in the neighbourhood of the reform's implementation. In reality the full impact of the reform will be smoothed over a number of cohorts. To understand the full dynamic impact on matching behaviour would entail looking at the reform in a general equilibrium, which is beyond the scope of this paper, but a promising avenue for future research.

Chapter 5

Unemployment and Domestic Violence: Theory and Evidence

5.1 Introduction

During each global recession of the past decades there have been recurrent suggestions in the media that domestic violence increases with unemployment. In 1993, for example, the British daily newspaper *The Independent* cited a senior police officer as saying of the increase in domestic violence:

> "With the problems in the country and unemployment being as high as it is and the associated financial problems, the pressures within family life are far greater. That must exacerbate the problems and, sadly, the police service is now picking up the pieces of that increase." (Andrew May, Assistant Chief Constable South Wales, The Independent, 9 March 1993)

In a 2008 interview for *The Guardian*, the Attorney General for England and Wales argued that domestic violence will spread as the recession deepens:

> "When families go through difficulties, if someone loses their job, or they have financial problems, it can escalate stress, and lead to alcohol or drug abuse. Quite often violence can flow from that." (Baroness Scotland of Asthal, The Guardian, 20 December 2008)

And in 2012, the executive director of a Washington-based law enforcement think-tank expressed his concerns about rising domestic violence rates in a *USA Today* article:

> "You are dealing with households in which people have lost jobs or are in fear of losing their jobs. That is an added stress that can push people to the breaking point." (Chuck Wexler, USA Today, 29 April 2012)

All these accounts are based on the same underlying logic and suggest that high unemployment may provide the "trigger point" for violent situations in the home. Yet, from a research perspective, it is far from clear whether unemployment is the overwhelming determinant of domestic violence that many commentators a priori expect it to be.[1] Indeed, a basic intra-household bargaining model would suggest that what really matters is the gender-profile of unemployment: an increase in male unemployment and/or a decrease in female unemployment should improve females' relative bargaining power, thereby reducing violence against women in much the same way as a decrease in the gender wage gap (Aizer, 2010).[2]

However, recent empirical evidence also points to factors such as emotional cues (Card and Dahl, 2011) and alcohol (Angelucci, 2008) as potential triggers of partner abuse. Such findings cast doubts on a theory that portrays partner abuse as intentional and rational acts that occur as part of Pareto efficient bargained outcomes. Nevertheless, even in settings where abuse is not an intentionally chosen action, as long as potential abusers have some influence over the likelihood of violent conflicts one would expect the economic logic of the bargaining power argument to carry over.

To verify this, we first present a novel model of partner abuse with gender-specific unemployment risk and where marriage allows couples to partially diversify income risk through consumption sharing. The model, which has parallels to the behavioural framework provided by Card and Dahl (2011), has the innovative feature that a female does not know her husband's "type" with regard to his predisposition to violence. For a given couple, the male partner may or may not have a violent predisposition, and his spouse infers his true nature from his behaviour. In equilibrium, a male with a violent predisposition can either reveal or conceal his type. When a male with a violent predisposition faces a high unemployment risk, he has an incentive to conceal his true nature by mimicking the behaviour of non-violent men. This follows since his spouse, given his low expected future earnings, would have a strong incentive to leave him if she were to learn his violent nature. As a consequence, higher male unemployment is associated with a lower risk of violence. Conversely, when a female faces a high unemployment risk, her low expected future earnings would make her less inclined to leave her partner even if she were to learn

[1] Specifically, we focus on violence against women perpetrated by their partners. While the term "domestic violence" generally also includes violence between other individuals within households, we will refer to partner violence and domestic violence interchangeably.

[2] The bargaining model provided by Aizer (2010) can be readily extended to include unemployment risk to derive this empirical prediction. Such an extension is available in Appendix D.4.

that he has a violent nature. Anticipating this, a male with a violent predisposition has weakened incentives to conceal his true nature, thus making higher female unemployment associated with a higher risk of violence. We therefore argue that, at a general level, a robust prediction from economic theory is that a high relative male unemployment rate should strengthen the relative economic position of females and lower partner violence against them.

Our empirical analysis combines high-quality individual-level data on intimate partner violence from the British Crime Survey (BCS) with local labour market data at the Police Force Area (PFA) level from the UK's Annual Population Survey (APS). Our basic empirical strategy exploits the substantial variation in the change in unemployment across PFAs, gender, and age-groups associated with the onset of the late-2000s recession. Our main specification links a woman's risk of being abused to the unemployment rates for females and males in her local area and age group. We first use basic probit regressions to estimate the effects of total and gender-specific unemployment rates on both physical and non-physical abuse. The structure of our data allows us to control for observable socioeconomic characteristics at the individual level as well as observable economic, institutional and demographic variables at the PFA level. In addition, we control for unobservable time-invariant area level characteristics and national trends in the incidence of abuse through the inclusion of area and time fixed effects. Finally, as our basic regressions suggest that unemployment matters for the incidence of abuse primarily because of the difference in unemployment rates by gender within areas and age groups, we instrument for the unemployment gender gap by exploiting differential trends in unemployment by industry and variation in initial local industry structure.

We find no evidence to support the common perception that domestic violence increases with the *overall unemployment rate*. This result parallels findings in previous studies suggesting near zero effects of total unemployment on domestic violence (Aizer, 2010; Iyengar, 2009). However, when we model the incidence of domestic violence as a function of *gender-specific unemployment rates*, as suggested by economic theory, we find that male and female unemployment have opposite-signed effects on domestic violence: while female unemployment increases the risk of domestic abuse, unemployment among males reduces it. The effects are also quantitatively important: the estimates imply that a 3.7 percentage point increase in male unemployment, as observed in England and Wales over the sample period, 2004 to 2011, causes a *decline* in the incidence of domestic abuse by up to 12%. Conversely, the 3.0 percentage point increase in female unemployment observed over the same period causes an *increase* in the incidence of domestic abuse by up to 10%. Thus, our results provide strong support for the predictions arising from the theory. We perform a battery of robustness checks on our data and find that our results are maintained across various alternative specifications. We further note that the relationship between gender-specific unemployment and partner abuse is unique to this type of crime: for the same group of respondents we do not find the same relationship to the personal experience of theft and general violence.

The paper contributes to a small but growing literature in economics on domestic violence. These studies can be divided into three broad categories. The first examines the relationship between the relative economic status of women and their exposure to domestic violence. Aizer (2010) specifies and tests a simple model where (some) males have preferences for violence and partners bargain over the level of abuse and the allocation of consumption in the household.[3] The key prediction of the model is that increasing a woman's relative wage increases her bargaining power and monotonically decreases the level of violence by improving her outside option. Consistent with this prediction, Aizer (2010) presents robust evidence that decreases in the gender wage gap reduce intimate partner violence against women.

The second type of study investigates the effects of public policy on domestic violence. Iyengar (2009) finds that mandatory arrest laws have the perverse effect of increasing intimate partner homicides. She suggests two potential channels for this: decreased reporting by victims and increased reprisal by abusers. Aizer and Dal Bó (2009) find that no-drop policies, which compel prosecutors to continue with prosecution even if a domestic violence victim expresses a desire to drop the charges against the abuser, result in an increase in reporting. Additionally, they find that no-drop policies also result in a decrease in the number of men murdered by intimates suggesting that some women in violent relationships move away from an extreme type of commitment device, i.e., murdering the abuser, when a less costly one, i.e., prosecuting the abuser, is offered.

The third type of study focuses more closely on male motives for violence. Card and Dahl (2011) argue that intimate partner violence represents expressive behaviour that is triggered by payoff-irrelevant emotional shocks. They test this hypothesis using data on police reports of family violence on Sundays during the professional football season in the US. Their result suggests that upset losses by the home team (i.e., losses in games that the home team was predicted to win) lead to a significant increase in police reports of at-home male-on-female intimate partner violence. Bloch and Rao (2002) argue that some males use violence to signal their dissatisfaction with their marriage and to extract more transfers from the wife's family. They test their model using data from three villages in India. Pollak (2004) presents a model in which partners' behaviour with respect to domestic violence is transmitted from parents to children.

[3] Earlier studies that have also employed a household bargaining approach to analyse domestic violence include Tauchen, Witte and Long (1991) and Farmer and Tiefenthaler (1997).

The remainder of the paper is organised as follows. Section 5.2 establishes the theoretical prediction, linking gender-specific unemployment risks to the incidence of domestic violence against women. Section 5.3 describes the data that we use. Section 5.4 outlines the methodology we employ to test the main ideas behind the model and presents the results. Section 5.5 concludes.

5.2 Theory

The main empirical hypothesis that we will take to the data is that the gender-profile of unemployment should matter for the incidence of domestic abuse. In order to verify the generality of this prediction—which follows naturally from a standard intra-household bargaining model—we present a novel theory of domestic violence in which abuse is not an intentionally chosen action and where asymmetric information occurs. Indeed, our model is the first economic theory to examine domestic violence in a setting where wives do not have perfect information about their husbands' types. The model is based on the premise that marriage is a non-market institution that can provide some degree of insurance against income risk. A key feature of our framework is that a male may or may not have a violent predisposition and that his female partner infers his type from his behaviour. In equilibrium, a male with a violent predisposition can either reveal or conceal his type, and his incentives for doing so depend on each partners' *future* earnings prospects as determined by their idiosyncratic unemployment risks and potential wages.

5.2.1 A Signaling Model with Forward-Looking Males

We consider a dynamic game of incomplete information involving two intimate partners: a husband (h) and a wife (w). The precise timing of the game is as follows:

1. Nature draws a type for the husband from a set of two possible types $\theta \in \{N, V\}$. Type V has a violent predisposition, while type N has an aversion towards violence. The probability that $\theta = V$ is denoted $\phi \in (0, 1)$.

2. The husband learns his type θ and chooses a behavioural effort from a binary set, $\epsilon \in \{0, 1\}$, which, along with his type, determines the probability that future conflictual interactions with his spouse escalate into violence. The probability of violence occurring is denoted by $\kappa(\theta, \epsilon) \in [0, 1]$. We assume that the behavioural effort $\epsilon = 1$ reduces the risk of violence and that a husband of type N is less prone to violence than a husband of type V. Hence $\kappa(\theta, 1) < \kappa(\theta, 0)$ for each $\theta \in \{N, V\}$ and $\kappa(N, \epsilon) < \kappa(V, \epsilon)$ for each $\epsilon \in \{0, 1\}$. Making the effort $\epsilon = 1$ costs the husband ξ (measured in utility units). Effort $\epsilon = 1$ can therefore be interpreted as a

costly action for the husband that reduces the likelihood of him "losing control" in a marital conflict situation. For example, he may voluntarily avoid criminogenic risk factors, such as excessive consumption of alcohol, or he may deliberately reduce his exposure to emotional cues (Card and Dahl, 2011).

3. The wife observes the husband's action ϵ (but not his type θ) and updates her beliefs about his type to $\hat{\phi}(\epsilon)$. Given her updated beliefs, she then decides whether to remain *married* or whether to get *divorced*, a decision we denote by $\chi = \{m, d\}$. If the wife decides to terminate the relationship, each partner i suffers a divorce cost $\alpha_i \geqslant 0$ (which may be emotional).

4. Nature decides on employment outcomes. Each partner i ($i = h, w$) is employed or unemployed with probabilities $1 - \pi_i$ and π_i, respectively. If employed, partner i earns income $y_i = \omega_i$. If unemployed, each individual has an income of $y_i = b$, which can be interpreted as an unemployment benefit.[4] We assume that $b < \omega_i$ for each partner i. If still married, each spouse obtains a monetary payoff that depends on total household income $(y_i + y_j)$. Formally, the monetary payoff of partner i is

$$u_i^m = \lambda^m + v(y_i + y_j), \tag{5.1}$$

where λ^m is a constant and where $v(\cdot)$ is an increasing, strictly concave and continuously differentiable function. If divorced, each partner's monetary payoff depends simply on his or her own income,

$$u_i^d = \lambda^d + v(y_i), \tag{5.2}$$

where λ^d is a constant which satisfies $\lambda^d \geqslant \lambda^m$. Note that our assumptions on relative payoffs are consistent with a fairly broad class of non-cooperative models of intra-household public good provision, the equilibria of which feature "local income pooling" but also inefficiency to due to "free-riding" (Warr, 1983; Bergstrom *et al.*, 1986; Browning et al., 2010).[5]

[4] The benefit income could be gender-specific, but we ignore this for notational simplicity.

[5] To see this, consider the simplest possible public good game in which the preferences of each individual i are represented by the utility function $u_i(x_i, G) = \ln(x_i) + \ln(G)$, where G is the amount of a "household good", which is a pure public good to married spouses. Private consumption by individual i is denoted by x_i. A married couple's total expenditure on the public good is the sum of individual contributions. These are given by g_i, and so $x_i = y_i - g_i$ and $G = g_h + g_w$. Each married individual chooses their contribution to the public good to maximise their utility, taking the contribution of their partner as given. Thus, individual i chooses g_i to maximise $\ln(y_i - g_i) + \ln(g_h + g_w)$. It is straightforward to show that both spouses will contribute to the household good for income shares satisfying $y_h/(y_h + y_w) \in (1/3, 2/3)$. In this case, $x_i^* = G^* = (y_h + y_w)/3$, i.e., there is "local income pooling" with household demands for all goods only depending on aggregate household income and not on individual incomes. Thus, the indirect

5. If still married, the couple encounters a conflict situation (e.g., heated disagreements) which escalates to violence with probability $\kappa(\theta, \epsilon)$. The wife suffers additive disutility $\delta_w > 0$ if violence occurs. The husband's disutility from violence is type-dependent, $\delta_N > 0$ for a husband of type N and $\delta_V = 0$ for a husband of type V.

We solve the model for a pure strategy perfect Bayesian equilibrium. Throughout, (ϵ', ϵ'') denotes that a husband of type V chooses ϵ' and a husband of type N chooses ϵ''. Similarly, (χ', χ'') indicates that the wife plays χ' following $\epsilon = 0$ and χ'' following $\epsilon = 1$.

5.2.2 Equilibrium

The wife rationally chooses whether or not to continue the marriage. Her expected payoff from getting divorced is given by:

$$D(\pi_w) = \mathbf{E}[u_w^d | \pi_w] - \alpha_w, \tag{5.3}$$

where

$$\mathbf{E}[u_w^d | \pi_w] = \lambda^d + (1 - \pi_w)v(\omega_w) + \pi_w v(b). \tag{5.4}$$

The expected value to the wife of remaining married depends not only on the wife's own unemployment risk, but also on the husband's unemployment probability and the perceived risk of domestic violence. Formally, the wife's expected payoff from remaining married is given by:

$$M(\pi_h, \pi_w, \epsilon, \hat{\phi}(\epsilon)) = \mathbf{E}[u_w^m | (\pi_h, \pi_w)] - \delta_w \left[(1 - \hat{\phi}(\epsilon))\kappa(N, \epsilon) + \hat{\phi}(\epsilon)\kappa(V, \epsilon) \right], \tag{5.5}$$

where

$$\begin{aligned} \mathbf{E}[u_w^m | (\pi_h, \pi_w)] = &\lambda^m + (1 - \pi_h)(1 - \pi_w)v(\omega_w + \omega_h)) + \pi_h \pi_w v(2b) \\ &+ \pi_h(1 - \pi_w)v(\omega_w + b) + \pi_w(1 - \pi_h)v(b + \omega_h). \end{aligned} \tag{5.6}$$

Note that the wife's expected utility from remaining married is decreasing in her perceived probability that the husband has a violent predisposition, $\hat{\phi}(\epsilon)$. The wife continues the partnership if and only if her expected value of remaining married exceeds the expected

utility function of individual i can be written as $u_i^m(x_i^*, G^*) = \lambda^m + v(y_h + y_w)$, where $\lambda^m = -2\ln(3)$ and $v(\cdot) = 2\ln(\cdot)$. A divorced individual chooses g_i to maximise $\ln(y_i - g_i) + \ln(g_i)$, and so the indirect utility function of a divorced individual can be written as $u_i^d(x_i^*, g_i^*) = \lambda^d + v(y_i)$, where $\lambda^d = -2\ln(2)$ and $v(\cdot) = 2\ln(\cdot)$. Notice that this simple example implies our assumption that $\lambda^d > \lambda^m$, which follows from the fact that married spouses "crowd out" each other's contributions to the public good, i.e., that their contributions are strategic substitutes.

value of getting divorced. The key assumptions of the model are as follows (for expositional convenience, we suppress the arguments of the functions):

A 1. *$M < D$ when $\pi_w = 0$, $\pi_h = 1$, $\epsilon = 0$ and $\hat{\phi} = 1$.*

A 2. *$M > D$ when $\pi_w = 1$, $\pi_h = 0$, $\epsilon = 0$ and $\hat{\phi} = 1$.*

A 3. *For any $(\pi_h, \pi_w) \in [0,1]^2$ and $\epsilon \in \{0,1\}$, $M > D$ when $\hat{\phi} = \phi$.*

The first two assumptions imply that the wife's tolerance of violence depends on her earnings prospects. To be more precise, suppose the wife observes the husband choosing $\epsilon = 0$. Assumption A1 ("not-take-it-if-employed") then says that if the wife will be *employed with certainty* and the husband will be *unemployed with certainty*, and she knows that the husband has a violent predisposition, then she will choose to divorce the husband. This may be interpreted as implying that economically independent women leave their abusive partners. On the other hand, assumption A2 ("accept-it-if-unemployed") implies that if the wife will be *unemployed with certainty* and the husband will be *employed with certainty*, and she knows that he has a violent predisposition, then she will not leave him. This captures the idea that women who are economically dependent on their abusers may be unable to leave them. Finally, assumption A3 ("stay-if-no-new-info") says that if the wife retains her prior beliefs, then she will continue the relationship irrespective of their unemployment probabilities and the husband's action. It is therefore consistent with wife accepting to be in a partnership with the husband in the first place.

In addition, we make the following two-part assumption:

A 4. *(i) $[\kappa(N,0) - \kappa(N,1)]\delta_N > \xi$, and (ii) $\alpha_h > \kappa(N,0)\delta_N$.*

Part (i) implies that a husband with an aversion towards violence values the reduction in violence associated with making the effort $\epsilon = 1$ more than its cost. Part (ii) is a sufficient condition to ensure that continued marriage is preferable to divorce for each type of husband $\theta \in \{N, V\}$ at any effort level $\epsilon \in \{0,1\}$. Thus, the husband has no incentive to choose his behavioural effort in a way that triggers a divorce.

Next we define $\hat{\pi}_w(\pi_h)$ as the unemployment probability for the wife at which she, conditional on having observed the husband choosing $\epsilon = 0$ and knowing that the husband has a violent predisposition, is indifferent between continued marriage and divorce. Formally, $\hat{\pi}_w(\pi_h)$ is implicitly defined through:

$$M(\pi_h, \hat{\pi}_w(\pi_h), 0, 1) = D(\hat{\pi}_w(\pi_h)). \tag{5.7}$$

Equation (5.7) may fail to have a solution in the unit interval. However, the following lemma tells us that it will do so for *some* values of π_h.

Lemma 1. *There exist two values, π_h' and π_h'', satisfying $0 \leq \pi_h' < \pi_h'' \leq 1$ such that (5.7) has a solution $\hat{\pi}_w(\pi_h) \in [0,1]$ for every $\pi_h \in [\pi_h', \pi_h'']$. Moreover, $\hat{\pi}_w(\pi_h)$ is differentiable at any $\pi_h \in (\pi_h', \pi_h'')$ with $\partial\hat{\pi}_w(\pi_h)/\partial\pi_h > 0$. In addition, $\partial\hat{\pi}_w(\pi_h)/\partial\omega_w > 0$ and $\partial\hat{\pi}_w(\pi_h)/\partial\omega_h < 0$.*

Proof. See Appendix D.1 □

Figure 5.1 illustrates a case where $\pi' > 0$ and $\pi_h'' < 1$. The locus $\hat{\pi}_w(\pi_h)$ partitions the set of possible unemployment risk profiles, $(\pi_h, \pi_w) \in [0,1]^2$, into two non-empty subsets or "regimes":

$$R_0 \equiv \left\{(\pi_h, \pi_w) \,|\pi_h \geq \pi_h''\right\} \cup \left\{(\pi_h, \pi_w) \,|\pi_w \leqslant \hat{\pi}_w(\pi_h)\right\}, \tag{5.8}$$

$$R_1 \equiv \left\{(\pi_h, \pi_w) \,|\pi_h < \pi_h'\right\} \cup \left\{(\pi_h, \pi_w) \,|\pi_w > \hat{\pi}_w(\pi_h)\right\}. \tag{5.9}$$

An increase in the husband's wage ω_h expands regime R_1 by shifting the locus $\hat{\pi}_w(\pi_h)$ downwards. In contrast, an increase in the wife's wage ω_w expands regime R_0 by shifting the locus upwards.

The following proposition shows that the nature of the game's equilibrium depends on which regime the couple's unemployment risk profile (π_h, π_w) falls within. Since signaling games are prone to equilibrium multiplicity, we focus on pure strategy equilibria that satisfy the commonly used Cho-Kreps "intuitive criterion" (Cho and Kreps, 1987).

Proposition 1. *In each regime there is a unique pure strategy perfect Bayesian equilibrium that satisfies the "intuitive criterion"*

(a) If $(\pi_h, \pi_w) \in R_0$, then

$$[(\epsilon', \epsilon'') = (1,1), (\chi', \chi'') = (d, m), \hat{\phi}(0) = 1, \hat{\phi}(1) = \phi]$$

is a "pooling" equilibrium.

(b) If $(\pi_h, \pi_w) \in R_1$, then

$$[(\epsilon', \epsilon'') = (0,1), (\chi', \chi'') = (m, m), \hat{\phi}(0) = 1, \hat{\phi}(1) = 0]$$

is a "separating" equilibrium.

Proof. See Appendix D.1 □

To see that this describes a perfect Bayesian equilibrium, consider each regime in turn, starting with R_0. Here a pooling equilibrium occurs where both types of husbands make the costly effort that reduces the risk of violence. A husband without a violent

Figure 5.1: Critical locus $\hat{\pi}_w(\pi_h)$ separating regime R_1 and regime R_0

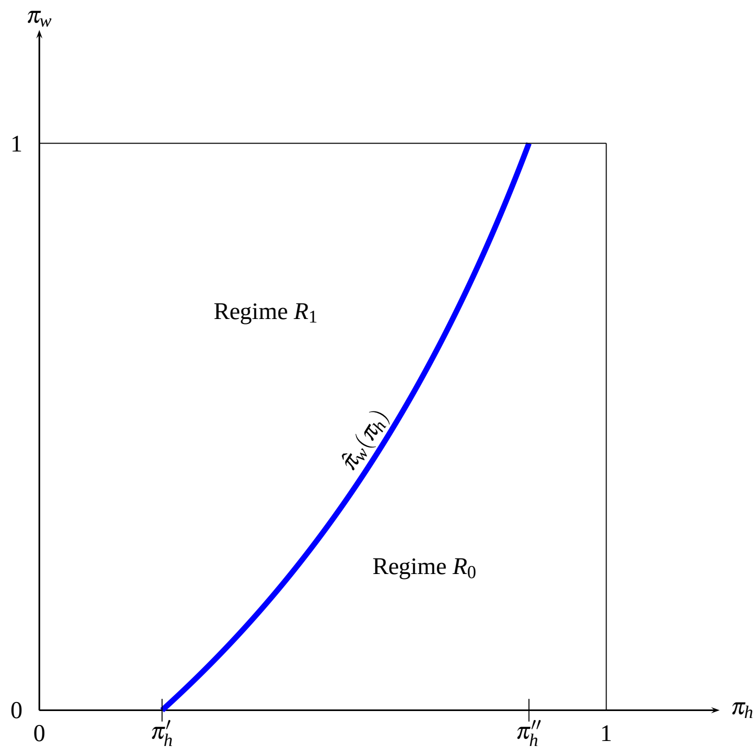

predisposition makes the effort since he values the reduction in the risk of violence that it generates more than the cost. A husband with a violent predisposition on the contrary makes the effort in order not to reveal his type as doing so would trigger a divorce. Central to the equilibrium are the wife's out-of-equilibrium beliefs and associated action: upon observing $\epsilon = 0$, the wife would conclude that the husband has a violent predisposition and would choose divorce.

Consider then regime R_1. In this case the husband knows that the wife is economically vulnerable and would not leave him even if she were to believe that he has a violent predisposition. A husband with a violent predisposition therefore has no incentives to make the costly effort that would reduce the risk of violence. A husband without a violent predisposition again values the reduction in the risk of violence more than the cost of making the effort. The wife's belief updating follows Bayes' rule and her continuing of the partnership with either type of husband is rational given her relatively weak earnings prospects.

5.2.3 Empirical Prediction

The above results form the basis of our empirical predictions: men with a violent predisposition may strategically mimic the behaviour of non-violent men, thus concealing their type, when facing relatively weak earnings prospects (Regime R_0) in the form of relatively high unemployment risk and relatively low wages. In contrast, when men face relatively strong earnings prospects (Regime R_1) they will be less inclined to conceal any violent predisposition they may have. Noting that the difference in the equilibrium probability of violence between Regime R_1 and R_0 is $\phi\left[\kappa\left(V,0\right)-\kappa\left(V,1\right)\right]>0$ we arrive at the following central empirical prediction:

Prediction 1.

- *A higher risk of male unemployment and lower wages for men are associated with a lower risk of domestic violence.*
- *A higher risk of female unemployment and lower wages for women are associated with a higher risk of domestic violence.*

Thus, we will build our empirical approach on the theoretical prediction that a woman's risk of being abused depends on gender-specific unemployment risks. In particular, in the empirical analysis we relate a woman's risk of experiencing domestic abuse to the local unemployment rates for males and females in her own age group.[6]

5.3 Data and Descriptive Statistics

5.3.1 Domestic Abuse Data from the British Crime Survey

We use data on the incidence of domestic abuse from the British Crime Survey (BCS). The BCS is a nationally representative repeated cross-sectional survey of people aged 16 and over, living in England and Wales, which asks the respondents about their attitudes towards and experiences of crime. The BCS employs two different methods of data collection with respect to domestic abuse. The first method, available from the survey's inception in 1981, is based on face-to-face interviews. However, the unwillingness of respondents to reveal instances of abuse to interviewers implies that this method significantly underestimates the true extent of domestic violence. To overcome such non-disclosure, a

[6] On a related note, Raphael and Winter-Ebmer (2001) find a negative association between female unemployment and state level rape rates which they suggest may be attributable to an increased number of encounters with potential perpetrators if an individual is working away from home. Rape perpetrated by a non-partner is, however, by definition not an outcome decided within a relationship, and therefore is not applicable to our theoretical framework.

Table 5.1: Demographic Characteristics of the BCS Sample

Variable	Mean	Std. Dev.	Variable	Mean	Std. Dev.
Age	38.93	11.67	Qual: Degree or above	0.236	0.425
Ethnicity: White	0.928	0.258	Qual: High Ed < Degree	0.137	0.344
Ethnicity: Asian	0.028	0.165	Qual: A level	0.150	0.357
Ethnicity: Black	0.023	0.150	Qual: GCSE grades A-C	0.237	0.426
Ethnicity: Other	0.021	0.143	Qual: Other	0.096	0.295
Religion: None	0.216	0.412	Qual: None	0.143	0.350
Religion: Christian	0.740	0.439	Single	0.355	0.479
Religion: Muslim	0.017	0.128	Married	0.455	0.498
Religion: Hindu	0.009	0.092	Separated	0.046	0.209
Religion: Sikh	0.004	0.060	Divorced	0.125	0.331
Religion: Jewish	0.003	0.057	Widowed	0.019	0.136
Religion: Buddhist	0.005	0.069	Cohabiting	0.120	0.325
Religion: Other	0.008	0.087	Long-standing illness	0.179	0.383
Number of children	0.493	0.896	Poor health	0.031	0.174
Children below 5 years	0.110	0.313			
Obs.			86,898		

self-completion module on interpersonal violence (IPV), which the respondents complete in private by answering questions on a laptop, was introduced.[7] We use BCS data for the survey years 2004/05 to 2010/11, covering interviews conducted between April 2004 and March 2011, and base our analysis on data on domestic violence from the self-completion IPV module.[8]

The BCS data has several key strengths as a source of data on domestic abuse. The IPV module in particular is unique in an international context, insofar that through self-completion the respondent does not need to provide answers directly to an interviewer. Furthermore, to reassure the respondent of privacy, the BCS randomly selects only one person per household who is exposed to the survey only once, implying that other household members including any partner will not know what questions the respondent has faced. In contrast the corresponding US survey, the National Crime Victimization Survey, administers the same set of questions to all household members every six months over a three year period, implying that the content of the questionnaire is common knowledge within the household.

Over our sample period, only 11 percent of those who report, in the IPV module, having been subjected to physical abuse by a partner also report being exposed to intra-

[7]The IPV module was first introduced in 1996. In 2001 it was used for a second time and the use of laptops was introduced. Since the 2004/05 survey the IPV module has been included on an annual basis, with a comparable set of questions.

[8]In the 2010-11 BCS survey, half of the sample were, in a trial, asked the same abuse questions, but in a simplified sequential format. For consistency we include in our sample only those respondents who were asked the abuse questions in the format consistent with the previous years' surveys.

Table 5.2: Categories of Domestic Abuse

behaviour	**Physical Abuse**	**Non-Physical Abuse**
Prevented from fair share of household money		x
Stopped from seeing friends and relatives		x
Repeatedly belittled you		x
Frightened you, by threatening to hurt you		x
Pushed you, held you down or slapped you	x	
Kicked, bit, or hit you	x	
Choked or tried to strangle you	x	
Threatened you with a weapon	x	
Threatened to kill you	x	
Used a weapon against you	x	
Used other force against you	x	

household abuse in the general interviewer-based part of the BCS survey. Similarly, only 48 and 50 percent report having mentioned the abuse to a medical staff and to the police respectively. Hence compared to alternative data from interviewer-based surveys, or data derived from police reports or hospital episodes statistics, the BCS IPV data is likely to provide substantially more comprehensive data on the incidence of domestic abuse. Furthermore, while police reports and hospital episode data can be used to measure incidence of (severe) domestic violence, such data generally cannot distinguish between multiple victims versus multiple events for the same victim. Finally, using micro-level data, obviously, allows us to control for individual level characteristics.

The BCS IPV module is answered by respondents aged 16 to 59, and we focus our analysis on intimate partner violence experienced by women.[9] Table 5.1 presents descriptive statistics of our sample.

In the IPV module respondents are presented with a list of behaviours that constitute domestic abuse and are asked to indicate which, if any, they have experienced in the 12 months prior to the interview. Table 5.2 presents this list of behaviours from which we, following the Home Office classification, construct two binary indicators of abuse. The first, *physical abuse*, is a dummy variable indicating whether the respondent had any type of physical force used against them by a current or former intimate partner. The second, *non-physical abuse*, indicates whether the respondent was threatened, exposed to controlling behaviours or deprived of the means needed for independence by a current or former partner.

In our sample, 3.0% of women report episodes of physical abuse in the past 12 months

[9] While the IPV module is also completed by male respondents, abuse against men is less common, generally less violent, and with no apparent connection to labour market conditions.

Figure 5.2: Incidence of Physical Abuse by Demographic Characteristics

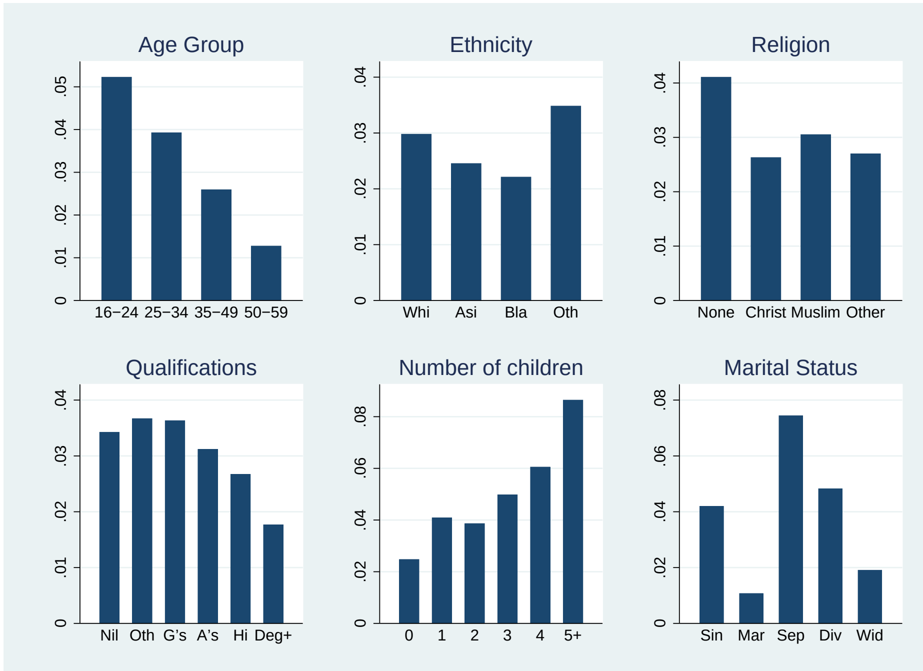

and 4.4% declare having experienced non-physical abuse.[10] Figure 5.2 illustrates the extent to which the incidence of physical abuse in particular varies with the demographic characteristics of the respondents. In general, exposure to physical abuse declines with age and with academic qualifications acquired after compulsory education. It varies relatively little with religion and ethnicity, but increases with the number of children.[11] With respect to marital status, it should be noted that this refers to the respondent's formal status at the time of the interview, which is hence observed *after* the 12 month period to which the abuse questions refer. The high reported rate of abuse among separated and divorce women therefore suggests a "reverse causality". The high rate of incidence among singles also emphasises the fact that "intimate partners" include current and past boyfriends.[12] Due to the highly endogenous nature of the respondent's current marital status we do not make use of this information except as a final sensitivity check on our estimates.[13] Figure 5.3 shows the trends in physical and non-physical abuse which, if anything, suggests

[10]The fraction of women reporting at least one of the two types of abuse was 5.7%.

[11]The relationship between physical violence and ethnicity is somewhat unexpected given the data from the US where blacks are typically found to have a higher incidence (Aizer, 2010).

[12]For respondents who are not currently married we also use a cohabitation dummy to indicate that the respondent is currently living with a partner. The incidence of abuse among currently cohabiting respondents is about double that of currently married respondents.

[13]The same applies to any information we have on the individual's current employment status. Hence we make no use of such information.

Figure 5.3: Trends in Domestic Abuse in England and Wales

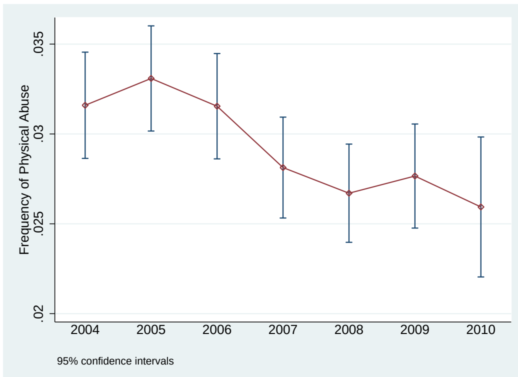

(a) Physical Abuse

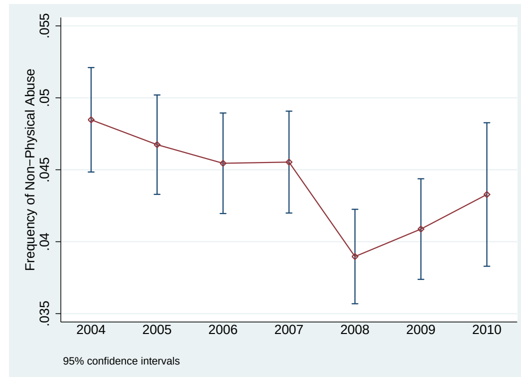

(b) Non-Physical Abuse

that the overall level of abuse is lower towards the end of our sample period than at the beginning. A corresponding decline has been observed over the same period in the most extreme form of violence against women: the rate of female homicide where the prime suspect is an intimate partner decreased by 6.3 percent between 2003-07 and 2007-11 (Smith et al., 2012).

5.3.2 Labour Market Data from the Annual Population Survey

We merge our individual-level data from the BCS with labour market data from the Annual Population Survey (APS). The APS combines the UK Labour Force Survey (LFS) with the English, Welsh and Scottish LFS boosts. Datasets are released quarterly, with each dataset containing 12 months of data. This means that we can, for each respondent in the BCS, using the known interview date, match the 12 month period to which the IPV questions refer to a closely corresponding 12 month period of labour market data.[14] Each respondent is matched to local labour market conditions corresponding to the Police Force Area (PFA) of residence, of which there are 42 in our data.[15]

[14] For instance, any respondent interviewed in the first three months of 2005 is matched to the labour market data for the calendar year 2004, whereas a BCS respondent interviewed between April and June in 2005 is matched to labour market data for the period April 2004 to March 2005 etc.

[15] There are 43 PFAs in England and Wales. However, the City of London PFA is a small police force which covers the "Square Mile" of the City of London. As this is a small area enclosed in the many times larger Metropolitan PFA we merge the two. This leaves us with 42 PFAs. They are Avon and Somerset, Bedfordshire, Cambridgeshire, Cheshire, Cleveland, Cumbria, Derbyshire, Devon and Cornwall, Dorset, Durham, Essex, Gloucestershire, Greater Manchester, Hampshire, Hertfordshire, Humberside, Kent, Lancashire, Leicestershire, Lincolnshire, City of London and Metropolitan Police District, Merseyside, Norfolk, Northamptonshire, Northumbria, North Yorkshire, Nottinghamshire, South Yorkshire, Staffordshire, Suffolk, Surrey, Sussex, Thames Valley, Warwickshire, West Mercia, West Midlands, West Yorkshire, Wiltshire, Dyfed-Powys, Gwent, North Wales and South Wales. The APS data is available in a finer geography, and is hence aggregated up to the PFA level.

Table 5.3: Summary Statistics for Local Unemployment Rates

Variable	**Mean**	**Std. Dev.**	**Min**	**Max**
Total unemployment	0.060	0.020	0.022	0.129
Unemployment by gender				
Male	0.064	0.023	0.022	0.149
Female	0.054	0.018	0.014	0.103
Unemployment by age group				
aged 16-24	0.150	0.045	0.029	0.283
aged 25-34	0.055	0.021	0.009	0.136
aged 35-49	0.039	0.016	0.010	0.104
aged 50-64	0.035	0.014	0.004	0.086

Notes: The table provides averages over the time-interval January 2003-December 2010 based on data from the APS which is provided in overlapping 12 month periods: January-December, April-March, July-June, October-September. Reported standard deviations and minimum and maximum values are over 1,218 PFA-period observations.

Our theory developed in the previous section stresses the role of male and female unemployment *risk* for the incidence of domestic violence. In the empirical analysis we will relate the incidence of domestic violence to the *observed* unemployment rates for the respondent's female and male peers, as defined by age group and geographical area. Hence we effectively interpret the observed unemployment rate not only as a measure of the direct incidence of unemployment, but also more broadly as an indicator for the perceived risk of unemployment. This interpretation is supported by the literature that documents workers' subjective unemployment expectations and relates it to the current level of unemployment. For instance for the US, Schmidt (1999) shows how workers' average beliefs about the likelihood of job loss in the next 12 months closely tracked the unemployment rate over the period 1977-96. The limited data that is available on unemployment expectations in the UK equally supports the notion that individual expectations of future unemployment risk are positively associated with the current unemployment rate. The British Social Attitudes (BSA) survey has, in selected years, asked respondents: (i) how "secure" they feel in their jobs, and (ii) whether they expect to see a change in the number of employees in their workplace. Both variables saw changes with the onset of the latest recession. In 2005, 78 percent of respondents reported feeling secure in their jobs; in 2009-2010, this figure had dropped to 73 percent. Similarly, while 16 percent of respondents reported expecting a reduction in the number of employees in the workplace in 2006-2007, this number had increased to 26 percent in 2009-2010.[16]

Table 5.3 presents basic descriptive statistics for local unemployment rates, broken

[16]Using data from the Skills Surveys, Campbell et al. (2007) document a similar fall in the average individual expectations of job loss between 1997 and 2001, a period of declining unemployment.

Figure 5.4: Gender-Specific Unemployment Rates and the Female-Male Unemployment Gap by Age Group in England and Wales, 2003 to 2011

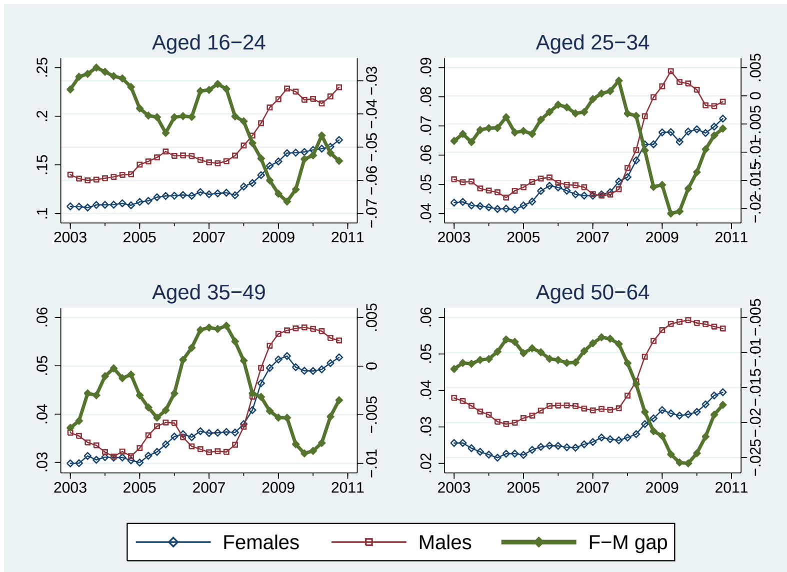

down by gender and age group.[17] Figure 5.4 shows that the increase in the rate of unemployment (left-hand scale) associated with the latest recession was far from uniform across gender and age groups. In particular, the impact of the recession is reflected more strongly in male than in female unemployment. As a consequence, we observe a widening of the female-male unemployment gap (right-hand scale) in the latter part of the sample period. In addition to local unemployment, we also use the APS to construct measures of mean hourly real wages.

Figure 5.5 contrasts the change over the sample period from 2004/05 to 2010/11 in the incidence of physical abuse with corresponding changes in male and female unemployment rates across the 42 PFAs. The figure highlights substantial spatial variation in the change in unemployment over the sample period. Moreover, the local changes in female unemployment are not obviously correlated with the corresponding local changes in male unemployment.[18] Inspection of the figure further suggests that several PFAs in which men were relatively more affected by unemployment increases (e.g., the North-East) saw relative decreases in the incidence of physical abuse. Indeed, if anything, the figure suggests a more positive association between relative increases in female unemployment and relative increases in abuse.

[17]The age grouping used in our analysis follows that conventionally used by the Office for National Statistics.

[18]Indeed, relevant to our identification strategy, less than forty percent of the variation in the gender unemployment gap is explained by area, period, and age group effects. This figure increases to just under fifty percent when age-group-period and age-group-area fixed effects are introduced.

Figure 5.5: Change in Female and Male Unemployment, in the Female-Male Unemployment Gap, and in Incidence of Physical Abuse across Police Force Areas in England and Wales, 2004 to 2011

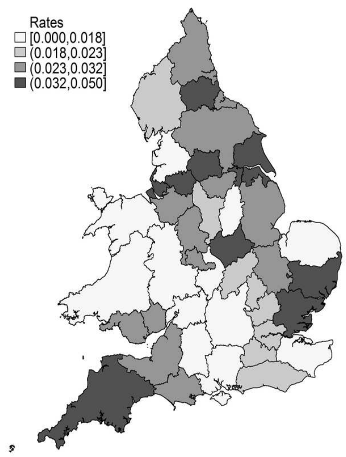

(a) Female unemployment

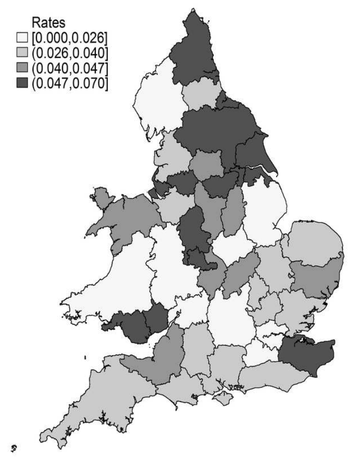

(b) Male unemployment

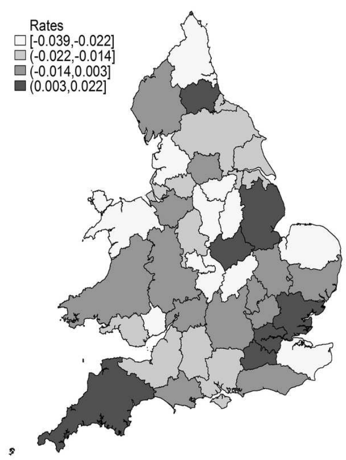

(c) Unemployment gap (F-M)

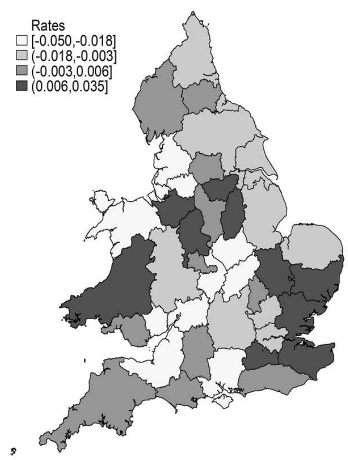

(d) Physical abuse

5.4 Empirical Specification and Results

5.4.1 Baseline Specification

This section presents our main analysis where we relate a female respondent's experience of domestic violence to the local level of unemployment. We focus in particular on the rates of female and male unemployment within the respondent's own age-group as these are likely to be the most relevant for the respondent's own unemployment risk as well as that of her (potential) partners. As the APS data is released quarterly, with each dataset containing 12 months of data, we define a "period" variable, denoted t, where a given period contains the particular APS release and BCS data from the following three months. Constructed in this way, our data stretches over 28 periods.[19]

As the outcome variables in our analysis are binary indicators of abuse, we estimate probit models. In particular, the basic model for the latent propensity for abuse against individual i in PFA j in period t and within age group g is given by

$$y^*_{ijtg} = \beta X_{ijtg} + \gamma^f UNEMPL^f_{jtg} + \gamma^m UNEMPL^m_{jtg} + \lambda_t + \alpha_j + \varepsilon_{ijtg} \tag{5.10}$$

where X_{ijtg} includes demographic controls at the individual level, $UNEMPL^f_{jtg}$ and $UNEMPL^m_{jtg}$ are the female and male unemployment rate in i's own age-group in police-force area j during period t, and ε_{ijtg} is a normally distributed random term.[20] The parameters λ_t and α_j are fixed effects for time-periods and police force areas respectively, and thus control for the aggregate trend in the outcome variable and for factors affecting abuse that vary across areas but are fixed over time. Thus, our basic model identifies the impact of gender-specific unemployment on domestic abuse from variation in trends across PFAs.

5.4.2 Baseline Results

Our basic results for the probability of being a victim of *physical abuse* are provided in Table 5.4.[21] Specification (1) gives the average marginal effect of the *total unemployment rate* within the own age group on the incidence of physical abuse. The estimated model includes a set of individual-level demographic controls: age measured in years and dummy variables indicating ethnicity, qualification level, and religious denomination, along with number of children and a dummy for the presence of at least one child under the age of five in the household. It further includes area- and time fixed-effects. The marginal effect is small and insignificant.[22] This result parallels findings in previous studies (Aizer, 2010; Iyengar, 2009) suggesting near zero effects of total unemployment on domestic violence.

Specification (2) reports the estimated average marginal effect of each *gender-specific unemployment rate* within the own age group. The marginal effect of female unemployment in the own age group is positive and statistically significant. The magnitude of the coefficient suggests that a 1 percentage point increase in the own-age female unemployment rate causes an increase in the likelihood of the respondent being a victim of physical abuse by 0.097 percentage points or little over 3% of the sample mean. We also see that

[19]See footnote 14 for further details.
[20]In Section 4.3 we further include area-level controls.
[21]Estimates from linear probability models are very similar.
[22]A (non-reported) regression on aggregate unemployment - across genders *and* age groups - is also not significant, but also has less precision due to low local variation from the national trend.

Table 5.4: Impact of Unemployment on Physical Abuse - Main Specification

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
Unemployment in own-age group	-0.026 (0.018)								0.008 (0.019)
Female unemployment in own-age group		0.097** (0.027)	0.093** (0.027)	0.102** (0.028)	0.094** (0.027)	0.083** (0.029)	0.102** (0.028)	0.091** (0.035)	
Male unemployment in own-age group		-0.090** (0.021)	-0.097** (0.022)	-0.081** (0.027)	-0.089** (0.021)	-0.094** (0.029)	-0.067* (0.029)	-0.084* (0.037)	
Female unemployment in other age groups			-0.013 (0.065)						
Male unemployment in other age groups			-0.047 (0.055)						
Female real wage in own-age group				0.005 (0.009)					
Male real wage in own-age group				-0.001 (0.006)					
Female-Male UE gap in own age group									0.094** (0.022)
Area and time fixed effects	yes	yes	yes	yes	yes	yes	yes	yes	yes
Linear age-in-years control	yes	yes	yes	yes	yes	no	no	no	yes
Age group fixed effects	no	no	no	no	no	yes	yes	yes	no
Age group * Period FEs	no	no	no	no	no	no	yes	yes	no
Age group * Areas FEs	no	no	no	no	no	no	no	yes	no
Other demographic controls	yes	yes	yes	yes	yes	yes	yes	yes	yes
Area-specific linear trends	no	no	no	no	yes	no	no	no	no
Obs.					86,731				

Notes: Standard errors clustered on police force area and age group in parentheses. "Other demographic controls" include dummies for ethnicity category, qualification level and religious denomination, number of children, and a dummy to indicate the presence of at least one child under the age of five in the household. ** Significant at 1%. * Significant at 5%.

the estimated average marginal effect of male unemployment is negative and statistically significant. The magnitude of the coefficient indicates that a 1 percentage point increase in male unemployment in the respondent's own age group causes a decline in the risk of physical abuse by 0.090 percentage points – again about 3% of the sample mean.

Controls for female and male unemployment within age groups other than the own are added in specification (3). We find that male and female unemployment within the own age group still have opposite-signed effects on the risk of physical abuse while unemployment in age groups other than the own appears to have little impact. Our theory suggests that potential wages of men and women might also matter for the incidence of abuse. Therefore, we add measures of local female and male mean hourly real wage rates within the own age group in specification (4). Controlling for wage-effects in this way leaves the marginal effects for male and female unemployment largely unchanged. The estimated wage effects are small and insignificant.[23] Specification (5) shows that our estimates are

[23]In fact, the coefficient have the "wrong" signs. In order to look further into this we obtained alternative measures of local wages from the Annual Survey of Hours and Earnings (ASHE) which is based on a one per cent sample of individuals from National Insurance records. Using this alternative data source, the

robust to the introduction of area-specific linear time trends.

Specifications (1) - (5) use the respondent's age-in-years as a control variable. This has the advantage of allowing us to use the exact information on the respondent's age. In contrast, our labour market variables are measured at the age group level. The next three specifications verify that our results are robust to alternate age controls. In specification (6), we replace the age-in-years variable with dummy variables indicating the respondent's age group, thus allowing for age group fixed effects along with the area and period fixed effects. In specification (7), we interact the age group dummies with the period dummies, thus allowing each age group to have a separate non-linear trend. In specification (8), we further interact the age group dummies with the area dummies, thus allowing each PFA to be associated with a separate fixed effect for each age group. Specification (8) is particularly restrictive as the fixed effects net out any changes in the unemployment rates across age groups over time as well as differences between age groups across areas. However, the point estimates remain of a similar magnitude to those in our preferred specification (2), although they lose some precision. This indicates that across all specifications our identification is primarily driven by age-area-time differences in unemployment rates.

An evident feature of the results in Table 5.4 is that the estimated effects of female and male unemployment are of very similar absolute magnitude, but of opposite sign. This suggests that what matters for the incidence of abuse is not the overall level of unemployment but rather the unemployment gender gap. Hence, in specification (9), we report the estimated marginal effect of the linear difference between the female and male unemployment rates within the own age group and of the total unemployment rate in the own age group. The estimated effect of the unemployment gender gap is noticeably strong whereas the estimated effect of the overall unemployment rate is not statistically significant. Specification (9) will also serve as the benchmark regression for our IV analysis below where we will also focus on the gender unemployment gap.[24]

Table 5.5 presents corresponding results for *non-physical abuse*. The estimated marginal effects for this alternative outcome variable are strikingly similar to those for physical abuse.

Table 5.6 breaks the estimated effect of the gender unemployment gap down by popula-

coefficient on wages have the expected sign, but remain statistically insignificant.

[24] As for our non-reported demographic control variables two factors stand out. Women with academic qualifications at A-level or above are less at risk of abuse. In contrast, there is a strong and significant positive correlation between the number of children and the incidence of abuse. To the extent that children reduce their mother's earnings capacity, this result is in line with our theoretical prediction. More generally, in our model any increase in the gains from marriage over divorce for the wife – obtaining from children or any other source – will be exploited by an abusive husband and make violence more likely. However, it is also possible that children are a cause of extra stress within a partnership, and that this provides a trigger for more violence. Appendix D.3 presents expanded versions of Tables 5.4 and 5.5 which include the coefficients on the demographic control variables.

Table 5.5: Impact of Unemployment on Non-Physical Abuse - Main Specification

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
Unemployment in own-age group	-0.012 (0.023)								0.021 (0.024)
Female unemployment in own-age group		0.102** (0.037)	0.108** (0.038)	0.110** (0.038)	0.104** (0.037)	0.078* (0.040)	0.093* (0.042)	0.087 (0.048)	
Male unemployment in own-age group		-0.081** (0.030)	-0.074* (0.032)	-0.061 (0.037)	-0.085** (0.031)	-0.090* (0.039)	-0.066 (0.039)	-0.077 (0.048)	
Female unemployment in other age groups			0.031 (0.080)						
Male unemployment in other age groups			0.035 (0.068)						
Female real wage in own-age group				-0.002 (0.010)					
Male real wage in own-age group				0.008 (0.007)					
Female-Male UE gap in own age group									0.093** (0.032)
Area and time fixed effects	yes	yes	yes	yes	yes	yes	yes	yes	yes
Linear age-in-years control	yes	yes	yes	yes	yes	no	no	no	yes
Age group fixed effects	no	no	no	no	no	yes	yes	yes	no
Age group * Period FEs	no	no	no	no	no	no	yes	yes	no
Age group * Areas FEs	no	no	no	no	no	no	no	yes	no
Other demographic controls	yes	yes	yes	yes	yes	yes	yes	yes	yes
Area-specific linear trends	no	no	no	no	yes	no	no	no	no
Obs.					86,731				

Notes: See notes to Table 5.4

tion subgroup in three dimensions. The top panel shows that the relationship is apparent for all bar the eldest age group. That nothing is found in the oldest age group is not entirely surprising given the low incidence of domestic violence reported in this age group.

The lower left panel in Table 5.6 splits the respondents into those with "high" educational attainment (A-level or above) versus those with "low" attainment (GCSE level or below).[25] One may argue that individuals' with lower qualifications are more at risk of unemployment and that, as a consequence, they may be more affected by gender unemployment gap in terms of the incidence of abuse.[26] While the point estimate is higher for low qualified women, the difference in the estimated effects is not statistically significant.

One may similarly argue that female unemployment is less relevant when the labour force participation (LFP) rate is relatively low. To consider this, we calculate the average female LFP over the sample period for each PFA-age group cell and partition the cells into those with above versus below median female LFP rate. Estimates by subgroup are reported in the lower right panel of Table 5.6. Again, while not statistically significantly

[25] The "high" qualifications are effectively those that require undertaking post-compulsory education.

[26] However, noting that the earnings drop associated with unemployment tends to be larger among individuals with higher qualifications, the effect could in principle go in the either direction.

Table 5.6: Impact of Unemployment Gender Gap on Abuse by Population Subgroup

	Age Group			
	16-24	25-34	35-49	50-59
	0.082**	0.122**	0.128*	-0.047
	(0.030)	(0.044)	(0.063)	(0.096)
	Own Qualification		*Female LFP in Cell*	
	"Low"	"High"	"Low"	"High"
	0.114*	0.089**	0.075**	0.137**
	(0.053)	(0.024)	(0.025)	(0.045)
Obs.	86,731			

Notes: The table reports average marginal effects from three probit estimations of the impact of the unemployment gender gap on physical abuse, with the same set of controls as in specification (2) in Table 4. ** Significant at 1%. * Significant at 5%.

different, the point estimates suggest that the effect of the gender unemployment gap on the incidence abuse is, if anything, stronger when the female LFP is higher.

The observed relationship between the gender-profile of unemployment and intimate partner violence can be expected to be particular to this outcome and not hold for general victim experience of crime. To verify this we replace our main outcome variables with other reported crime outcomes. The BCS respondents are asked whether, over the past 12 months, they have experienced theft from their person or been a victim of a violent assault.[27] The results of this analysis are displayed in Table 5.7. For both theft and violence we find, in line with the literature (Raphael and Winter-Ebmer, 2001; Öster and Agell, 2008), that the probability of reporting having been a victim of crime increases with total unemployment. Moreover, unlike domestic abuse where there can be expected to be a direct power relationship between the perpetrator and the victim, these outcomes if anything increase with both the male and the female rate of unemployment.

To summarise, consistent with the literature, we find no evidence to support the view that total unemployment increases domestic abuse. Instead, our results suggest that

[27]In both outcomes the victim is present at the time of the crime, so gender is readily identifiable. In the case of theft, as this crime is mainly an opportunist event, the gender of either victim or perpetrator should however play only a minor role. Considering violence, one might expect that in cases of affray the victim and assailant often share the same gender. The exact questions answered by the respondents were: "Was anything you were carrying stolen out of your hands or from your pockets or from a bag or case?" and "Has anyone, including people you know well, *deliberately* hit you with their fists or with a weapon of any sort or kicked you or used force or violence in any other way?".

Table 5.7: Impact of Unemployment on Experience of Crime

Specification	**Theft from Person**	**Theft from Person**	**Violence against Person**	**Violence against Person**
Total unemployment in own-age group	0.099** (0.017)		0.035* (0.015)	
Female unemployment in own-age group		0.042 (0.028)		0.039 (0.028)
Male unemployment in own-age group		0.056** (0.020)		0.004 (0.021)
Area and time fixed effects	yes	yes	yes	yes
Linear age-in-years control	yes	yes	yes	yes
Other demographic controls	yes	yes	yes	yes
Observations	86,725	86,725	86,726	86,726

Notes: See notes to Table 4.

male and female unemployment have distinct impacts on the incidence of domestic abuse: increases in male unemployment are associated with declines in domestic abuse while increases in female unemployment have the opposite effect. These findings are consistent with economic theory. The magnitude of the estimated relationships imply (a) that a 3.7 percentage point increase in male unemployment, as observed in England and Wales between 2004 and 2011, causes a *decline* in the incidence of domestic abuse of between 10.1% and 12.1%, and (b) that the 3.0 percentage point increase in female unemployment over the sample period causes an *increase* in the incidence of domestic abuse of between 9.1% and 10.3%.

5.4.3 Extended Results: Area Level Controls

Our estimates in the previous section would be biased if there were omitted variables that are correlated with local unemployment and that affect the incidence of domestic abuse. For example, a positive effect of unemployment on crime in general may trigger a response by the criminal justice system, such as increased police efforts or higher incarceration rates. If the response by the criminal justice system reduces domestic abuse by increasing deterrence, omitting controls related to the general level of criminal activity and the judiciary biases the estimated effect of unemployment on domestic abuse. Similarly, assuming that the consumption of alcohol and drugs is correlated with unemployment and also affects domestic abuse, omitting these factors from the regression again biases the estimates.[28] Additionally, selective migration might confound our estimates. For example, employment-driven migration of low-skilled men from areas with high local unemployment

[28] The association between business cycles and alcohol consumption is not clear cut. For instance, Dee (2001) notes that average drinking is generally pro-cyclical, but finds that binge-drinking is counter-cyclical.

Table 5.8

Impact of Unemployment on Physical Abuse and Non-Physical Abuse - Additional Controls

Specification	**(2)**	**(10)**	**(11)**	**(12)**	**(13)**	**(14)**	**(15)**
				(a) *Physical Abuse*			
Female unemployment in own-age group	0.097**	0.096**	0.102**	0.087**	0.097**	0.107**	0.092**
	(0.027)	(0.027)	(0.028)	(0.027)	(0.027)	(0.028)	(0.026)
Male unemployment in own-age group	-0.090**	-0.088**	-0.107**	-0.086**	-0.089**	-0.069**	-0.109**
	(0.021)	(0.021)	(0.021)	(0.025)	(0.021)	(0.026)	(0.021)
				(b) *Non-Physical Abuse*			
Female unemployment in own-age group	0.102**	0.101**	0.105**	0.091*	0.104**	0.109**	0.092*
	(0.037)	(0.038)	(0.038)	(0.039)	(0.037)	(0.039)	(0.037)
Male unemployment in own-age group	-0.081**	-0.080**	-0.090**	-0.077*	-0.083**	-0.073*	-0.104**
	(0.030)	(0.031)	(0.031)	(0.034)	(0.031)	(0.037)	(0.030)
Local area crime-related controls	no	yes	no	no	no	no	no
Local area drugs and alcohol	no	no	yes	no	no	no	no
Local area qualifications distribution	no	no	no	yes	no	no	no
Selective migration	no	no	no	no	yes	no	no
Unemployment in neighbouring areas	no	no	no	no	no	yes	no
Health and marital status	no	no	no	no	no	no	yes
Observations	86,731	86,731	80,011	86,731	86,731	86,731	86,674

Notes: Standard errors clustered on police force area and age group in parentheses. All specifications include area and time fixed effects, linear age-in-years control and other demographic controls (see notes to Table 4). Local area crime related-controls include police force manpower per 10,000 capita, violent and non-violent crimes per 10,000 capita, and average time from charge to magistrate court appearance. Local area drugs and alcohol includes the number of arrests for drugs possession per 10,000 capita and the number of alcohol-related hospitalisations per 10,000 capita. Selective migration includes the number of in- and out-migrants as a percentage of the PFA population in the respondent's own-age and gender group. For a detailed description of controls used in this section, see Appendix D.2. ** Significant at 1%. * Significant at 5%.

to areas with low local unemployment creates a downward bias (due to "compositional effects") if low-skilled males have a higher propensity to abuse their partners than high-skilled males. To mitigate such omitted-variables bias, we now control extensively for observable institutional and demographic covariates at the police-force area-level.

The results for *physical abuse* are shown in panel (a) of Table 5.8. Specification (2) repeats our basic specification from Table 5.4 for convenience. In specification (10), we add a set of controls that capture the general level of criminal activity and the potential response by the criminal justice system to it. In particular, we include per capita measures of violent and non-violent crimes. We include per capita measures of police force manpower and a proxy for the "efficiency" of the criminal justice system: the average time from charge to magistrate court appearance. Overall, the inclusion of these crime-related controls leaves our key estimates unchanged. This suggests that variation in overall crime rates and policing and criminal justice efforts do not confound our estimated effects of unemployment on domestic abuse.

Specification (11) includes a measure of the hospitalisation rate for alcohol-related conditions as well as a per capita measure of drugs possession.[29] Adjusting for the cyclical consumption of criminogenic commodities in this way does not alter our main finding that male and female unemployment have opposite-signed effects on the incidence of physical abuse. In specification (12), we account for the possibility of skill-selective migration by including the qualification distribution in the respondent's own-age group. Specification (13) controls directly for area-level migration by including the number of in- and out-migrants as a percentage of the PFA population in the respondent's own-age group. In each case, the estimated marginal effects of gender-specific unemployment remain largely unaffected.

The two remaining specifications provide additional robustness checks. Specification (14) shows that our results are robust to the introduction of controls for the average own-age group female and male unemployment rates in neighbouring police-force areas. Specification (15) shows that our main findings remain intact also when we include controls that capture a respondent's marital and health status (measured at the time of the interview and hence after the period to which the abuse information pertains).

Panel (b) of Table 5.8 provides the corresponding extended results for *non-physical abuse*. Again, the general conclusion is that the estimated effects of unemployment by gender are robust to the inclusion of further controls. The results presented in this section thus suggest that our initial finding that female unemployment increases domestic abuse while male unemployment reduces it is robust to including a wide variety of observable institutional and demographic covariates at the PFA level.

5.4.4 Instrumental Variables Estimation

The analysis so far has treated the local unemployment variables as exogenous regressors. Concerns about potential omitted variables motivated our use of additional regressors in Section 5.4.3. However, this may not have entirely solved the potential issue of omitted variables and would not address any potential problem of simultaneity. Solving these problems requires constructing measures of local labour market conditions that do not reflect characteristics of female and male workers, which could be affected by violence itself, or unobservables that might be correlated with violence. Hence as a final robustness check, we also consider an instrumental variables approach. Building on the work of Bartik (1991) and Blanchard and Katz (1992), we interact the initial local industry composition of employment with the corresponding national industry-specific trends in unemployment.

[29] Information on hospitalisation rates for alcohol-related conditions in particular is only available for England. This accounts for the drop in the number of observations in this particular specification.

Specifically, we use APS data on local PFA industry composition by gender and age group at baseline, defined as the calendar year 2003, which we combine with APS data on industry unemployment rates by gender and age group at the national level over the sample period.[30] For each PFA, gender, age-group and time period we construct an industry-predicted unemployment rate as follows,

$$\widehat{UNEMPL}^{h}_{jtg} = \sum_{k} \psi^{h}_{jgk} UNEMPL^{h}_{ktg}, \tag{5.11}$$

where ψ^{h}_{jgk} is the share of industry k among employed individuals of gender h and age group g in PFA j at baseline, and where $UNEMPL^{h}_{ktg}$ is the unemployment rate, at the national level, in industry k among individuals of gender h and age group g in time period t. Hence (5.11) is a weighted average of the national industry-specific unemployment rates where the weights reflect the baseline local industry composition in the relevant gender and age group. The weights are thus fixed over time and do not reflect local sorting into industries over the sample period.

Our approach draws on recent work by Albanesi and Sahin (2013) who, using US data, show how the gender gap in unemployment tends to vary over the business cycle. In particular, they find that unemployment rises more for men than for women during recessions, and also decreases more for men in subsequent recoveries. The authors also explore the role played by gender differences in industry structure. Specifically with respect to the recession in the late 2000s, Albanesi and Sahin show how gender differences in industry composition explain around half of the difference in the observed unemployment growth. Based on this observation, and on our previous finding that unemployment appears to matter for the incidence of domestic abuse only in the form of the unemployment gender gap, our IV analysis will be focused on estimating models where the incidence of domestic violence is related to the female-male unemployment gender gap. We instrument for the actual gender gap using the corresponding industry-predicted gender gap in unemployment.

Table 5.9 presents the results for two different specifications, each estimated as basic probit and as IV probit model. Specification (1) in Table 5.9 includes the same controls as in specification (2) in Table 5.4. Hence the difference is that here we include the

[30] Eight industries are used in the analysis based on a condensed version of the UK Standard Industrial Classification of Economic Activities, SIC(2007): "Agriculture, forestry, fishing, mining, energy and water supply", "Manufacturing", "Construction", "Wholesale, retail & repair of motor vehicles, accommodation and food services", "Transport and storage, Information and communication", "Financial and insurance activities, Real estate activities, Professional, scientific & technical activities, Administrative & support services", "Public admin and defence, social security, education, human health & social work activities", "Other services". The "industry unemployment rate" is defined as the unemployed by industry of last job as percentage of economically active by industry.

Table 5.9: Impact of Unemployment on Physical Abuse - Instrumental Variables Estimation

Specification	**(1a)** Probit	**(1b)** IV Probit	**(2a)** Probit	**(2b)** IV Probit
	(a) *Gender Unemployment Gap in Own Age Group*			
Predicted unemployment gender gap in own age group		1.733** (0.106)		1.723** (0.102)
	(b) *Physical Abuse*			
Gender unemployment gap in own age group	0.090** (0.021)	0.104* (0.049)	0.089** (0.021)	0.105* (0.049)
	(c) *Non-Physical Abuse*			
Gender Unemployment gap in own age group	0.081** (0.031)	0.083 (0.062)	0.084** (0.031)	0.081 (0.063)
Area and time fixed effects	yes	yes	yes	yes
Linear age-in-years control	yes	yes	yes	yes
Other demographic controls	yes	yes	yes	yes
Area-specific linear time trends	no	no	yes	yes
Observations			86,731	

Notes: See notes to Table 4.

unemployment rates in the own age group in the form of the gender gap rather than as levels. Specification (2) in Table 5.9 includes the same controls as in specification (5) in Table 5.4. The probit estimated average marginal effects of the gender unemployment gap on physical and non-physical abuse reported in columns (1a) and (2a) are naturally in line with the corresponding estimates in Tables 5.4 and 5.5.

Turning to the IV probit estimates, panel (a) of Table 5.9 confirms that our instrument is indeed a strong and relevant predictor of the gender unemployment gap in the own age group. More precisely, the estimates show that the actual variation in gender unemployment gap trends across PFAs and age groups is strongly positively related to the corresponding variation in the unemployment gap trends predicted using local variation in industry structure at baseline.

The IV probit estimated average marginal effects of the gender unemployment gap on the incidence of domestic abuse are reported in columns (1b) and (2b). For physical abuse we find that, for both specifications, the IV estimated marginal effects are slightly larger than, but not statistically significantly different from, the corresponding probit estimated effects. Each estimated marginal effect is also statistically significant. For non-physical abuse, the IV probit estimated average marginal effects of the gender unemployment gap are also very similar to the basic probit estimated effects. However, due to lower precision, they are not statistically significant. Overall, we view our IV estimates as evidence that our basic probit estimates do not exaggerate the impact of unemployment on domestic abuse.

5.5 Concluding Comments

This paper has examined the effect of unemployment in England and Wales on partner abuse against women. The geographical variation in unemployment in these countries induced by the Great Recession provides an interesting context in which to look at domestic abuse. Our empirical approach was motivated by a theoretical model in which partnership provides insurance against unemployment risk through the pooling of resources. The key theoretical result is that an increased risk of male unemployment lowers the incidence of intimate partner violence, while an increased risk of female unemployment leads to a higher rate of domestic abuse. We have demonstrated that this prediction accords well with evidence from the British Crime Survey matched to geographically disaggregated labour market data. In particular, our empirical results suggest that a 1 percentage point increase in the male unemployment rate causes a *decline* in the incidence of physical abuse against women of around 3 percent, while a corresponding increase in the female unemployment rate has the opposite effect. Moreover, our results also rationalise findings in previous studies of near zero effects of the *overall* rate of unemployment on domestic violence.

Overall, our theoretical model and empirical results contrast the conventional wisdom that male unemployment in particular is a key determinant of domestic violence. Quite the contrary, latent abusive males who are in fear of losing their jobs or who have lost their jobs may rationally abstain from abusive behaviours, as they have an economic incentive to avoid divorce and the associated loss of spousal insurance. However, when women are at a high risk of unemployment, their economic dependency on their spouses may prevent them from leaving their partners. This in turn might prompt male partners with a predisposition for violence to reveal their abusive tendencies. Thus, high female unemployment leads to an elevated risk of intimate partner violence. From a policy perspective, it is therefore conceivable that policies designed to enhance women's employment security could prove an important contributor to domestic violence reduction.

Appendix A

Appendices

A Appendix to Chapter 2

Figure A.1: Sharp RDD for all outcome variables

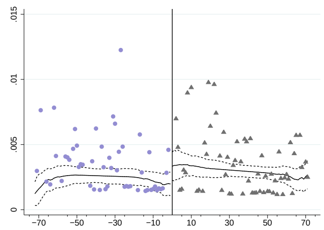

(a) Mother at age 15

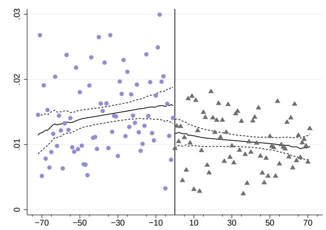

(b) Mother at age 16

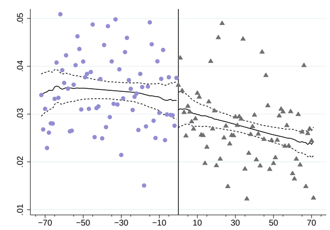

(c) Mother at age 17

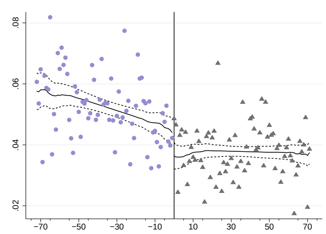

(d) Mother at age 18

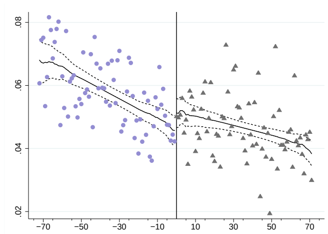

(e) Mother at age 19

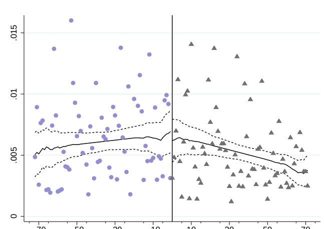

(f) Mother by age 16

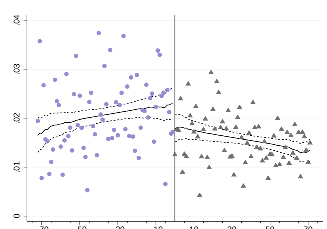

(g) Mother by age 17

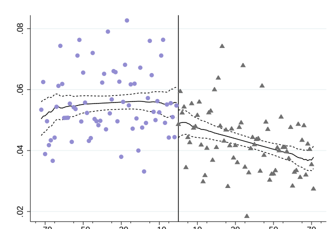

(h) Mother by age 18

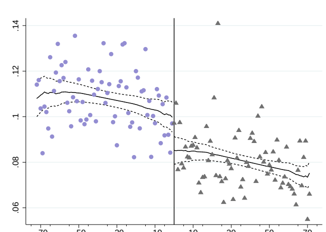

(i) Mother by age 19

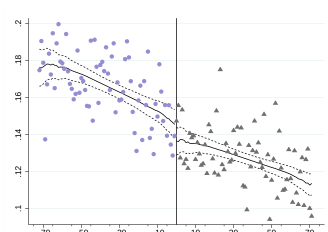

(j) Mother by age 20

Notes: The graphs display local-linear polynomial smooths, as described in Section 2.4.1, using a bandwidth of 48 months and a rectangular kernel, for the probability of becoming a mother at age 15 - age 19 (graphs a - e) and by age 16 - by age 20 (graphs f - j). The horizontal axis measures the distance, in months, of individuals' births to the RoSLA cutoff. The scatterplot indicates the proportions of mothers in each month-bin. The dashed lines are 95% confidence intervals of the local polynomial.

Figure A.2: Cross-Validation functions

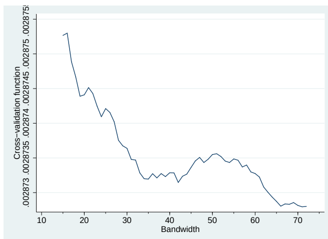

(a) Mother at age 15

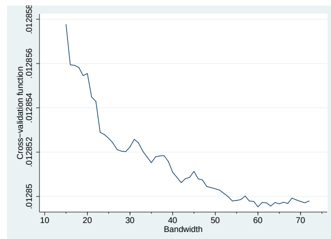

(b) Mother at age 16

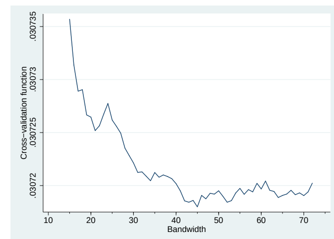

(c) Mother at age 17

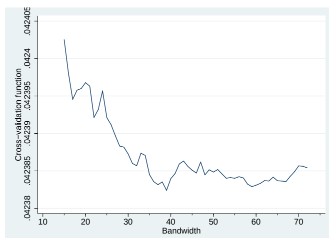

(d) Mother at age 18

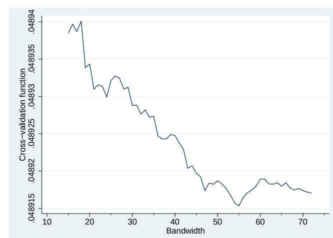

(e) Mother at age 19

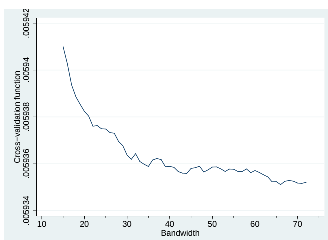

(f) Mother by age 16

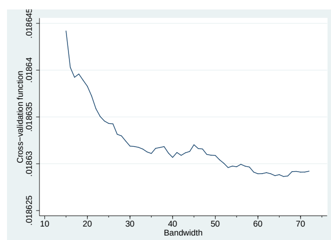

(g) Mother by age 17

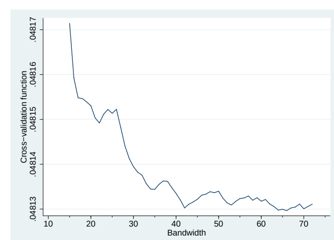

(h) Mother by age 18

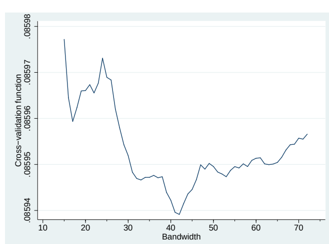

(i) Mother by age 19

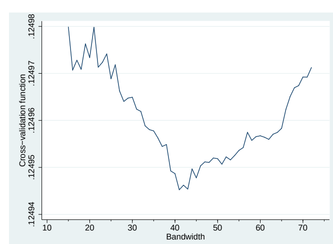

(j) Mother by age 20

Notes: The graphs display the cross-validation function for each of the outcome variables over the range of bandwidths following the procedure described in Section 2.4.3. The optimal bandwidth is defined as the value of the bandwidth, h, that minimizes the cross-validation function, $CV_Y(h)$, which is computed as the mean square difference of the predicted value to the true value of Y

Figure A.3: Sensitivity of Estimates to bandwidth choice

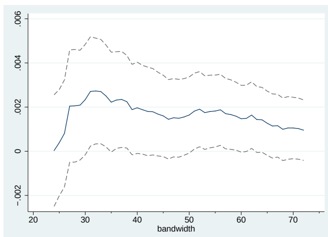

(a) Mother at age 15

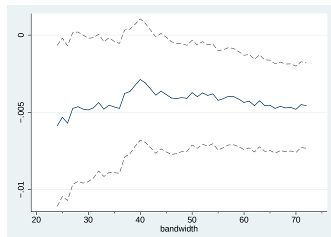

(b) Mother at age 16

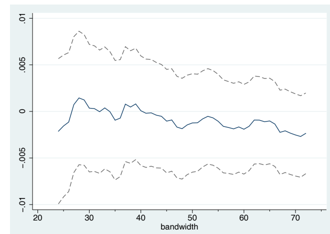

(c) Mother at age 17

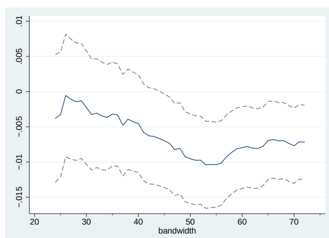

(d) Mother at age 18

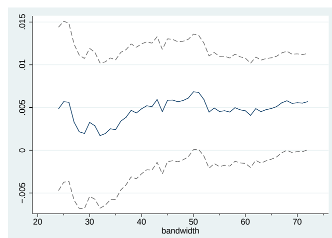

(e) Mother at age 19

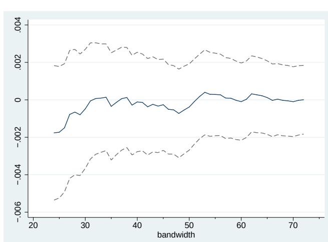

(f) Mother by age 16

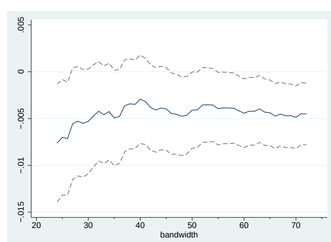

(g) Mother by age 17

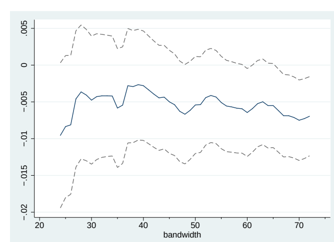

(h) Mother by age 18

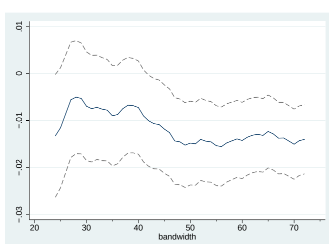

(i) Mother by age 19

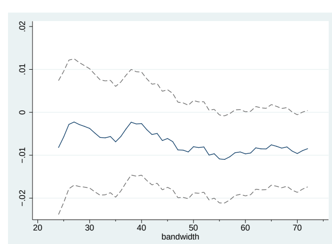

(j) Mother by age 20

Notes: The graphs display the magnitude of the estimates, along with the 95% confidence interval, over different bandwidths based on the parametric regression discontinuity design as described in Section 2.4.2.

B Appendix to Chapter 3

Table B.1: EMA Income Thresholds

Year	Median Income	Lower Threshold	% relative to median	Higher Threshold	% relative to median
1999/00	15400	13000	84	30000	195
2000/01	16100	13000	81	30000	186
2001/02	17100	13000	76	30000	175
2002/03	17800	13000	73	30000	169
2003/04	18200	13000	71	30000	165
2004/05	18900	19630	104	30000	159
2005/06	19400	20270	104	30000	155
2006/07	20400	20817	102	30810	151
2007/08	21200	20817	98	30810	145
2008/09	21000	20817	99	30810	147
2009/10	22200	20817	94	30810	139
2010/11	22500	20817	93	30810	137

Table B.2: Pilot period - participation in post-compulsory education

	Year 1 Participation			Year 2 Participation			Year 1&2 Participation		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: All Areas									
EMA	3.860***	3.731***	3.648***	2.973***	2.558***	2.510***	3.664***	2.945***	2.848***
	(0.486)	(0.436)	(0.446)	(0.397)	(0.385)	(0.390)	(0.355)	(0.343)	(0.341)
N	988	986	986	992	990	990	989	987	987
Panel B: Evaluation Areas									
EMA	3.361***	3.943***	3.516***	2.716***	1.835**	0.857	1.422*	1.320*	0.431
	(1.243)	(1.035)	(0.990)	(0.931)	(0.866)	(0.874)	(0.821)	(0.772)	(0.703)
N	156	156	156	156	156	156	156	156	156
Area controls	no	yes	yes	no	yes	yes	no	yes	yes
Region trends	no	no	yes	no	no	yes	no	no	yes

Notes: The table reports estimation of equation (3.5) as described in Section 3.4.2.1 during the pilot period of the EMA (Sept 1999 - Aug 2004). Panel A includes all non-pilot local authorities as the control group and are therefore comparable to the analysis of Chowdry et al. (2008), whereas Panel B is restricted to only those areas used in the final quantitative evaluation of the EMA pilots, and may be compared with the estimates of Middleton et al. (2005) (See footnote 3 for details of the treated and control areas). All specifications include area fixed effects and a quadratic aggregate year trend, and are weighted according to area population size. The vector of area controls includes the gender ratio, gender specific unemployment rates and real hourly wage, as well as area-level UK nationality and ethnic population proportions. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, ***$p < 0.01$

Figure B.1: EMA Timeline

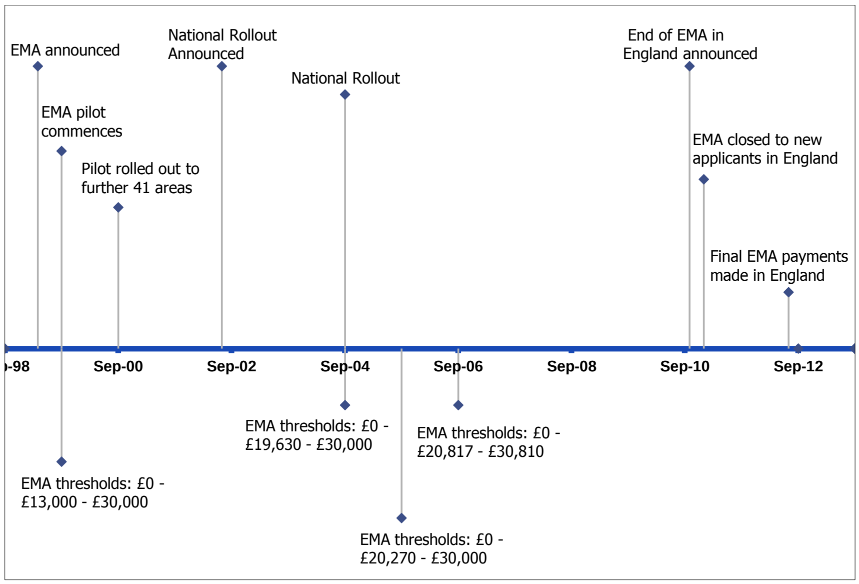

Figure B.2: Fertility Measures - 2007/08

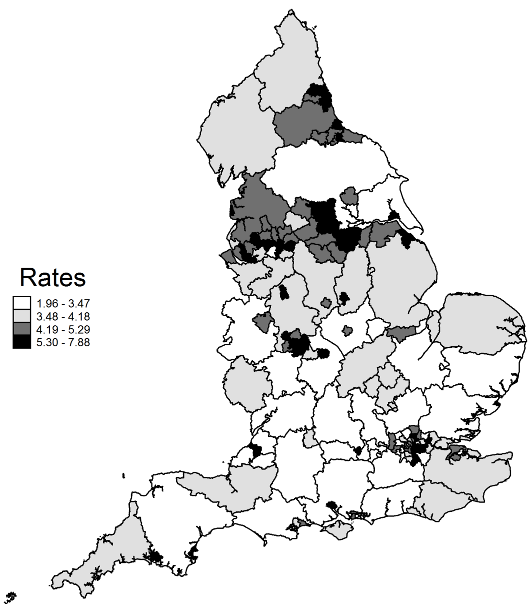

(a) Conceptions

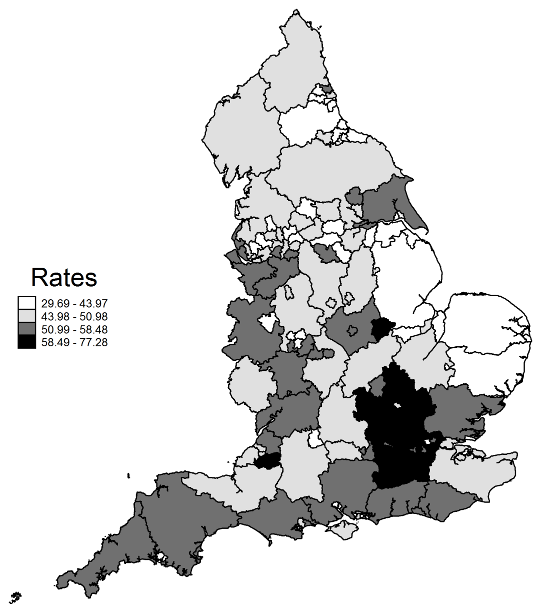

(b) Abortions

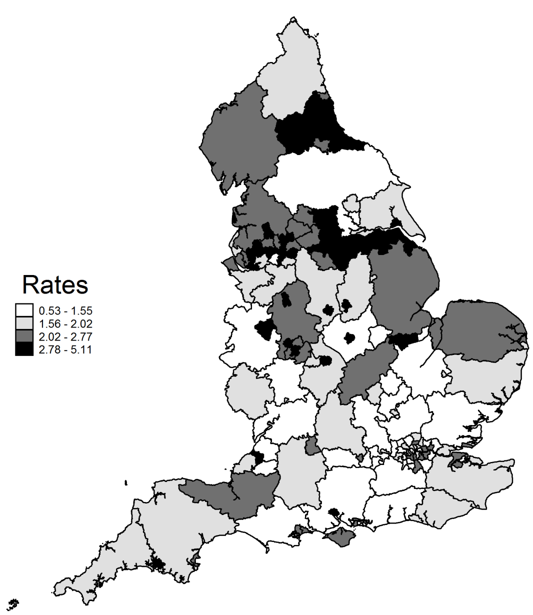

(c) Maternities

Notes: The maps display underage fertility rates for the academic year 2007/08. Conception rates is calculated as the number of conceptions to individuals under age 18 divided by the the area female 15-17 population. Map b) indicates the percentage of underage conceptions terminated; c) shows the underage maternity rate, defined by births to individuals who conceived before age 18.

Figure B.3: Crime Measures - 2007/08

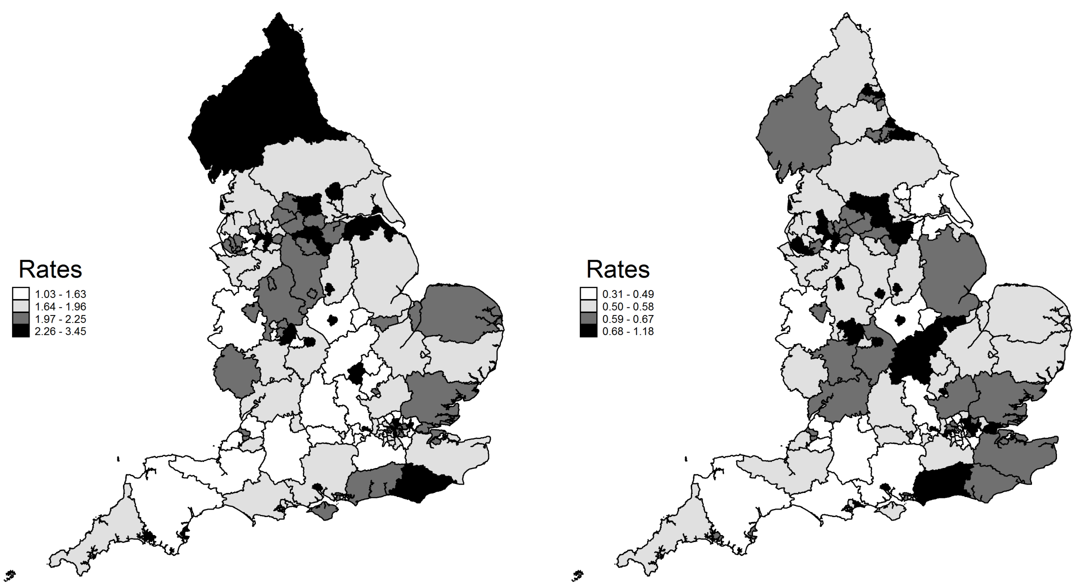

(a) 10-17 (b) Adults

Notes: The maps display the proportion of individuals in each area who are first time entrants to the Criminal Justice System in academic year 2007/08.

C Appendix to Chapter 4

C.1 Figures and Tables

Figure C.1: Marriage Age

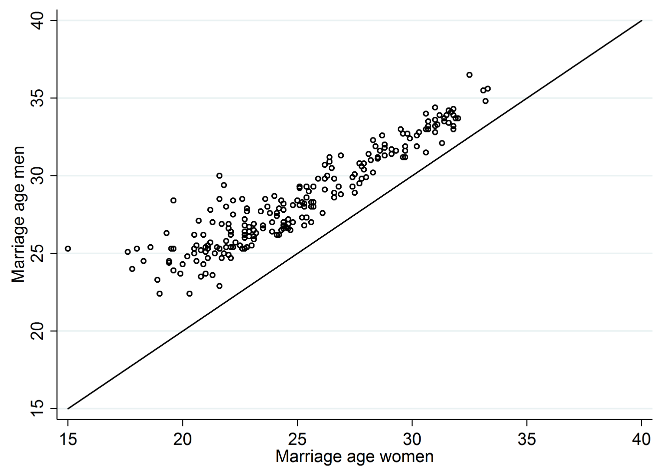

Notes: The graph displays singulate mean ages at marriage for men and women which are obtained from marital statuses by age for each of 218 countries where such data exists. Source: UN

Figure C.2: Marital age gap (husband-wife) by academic cohort

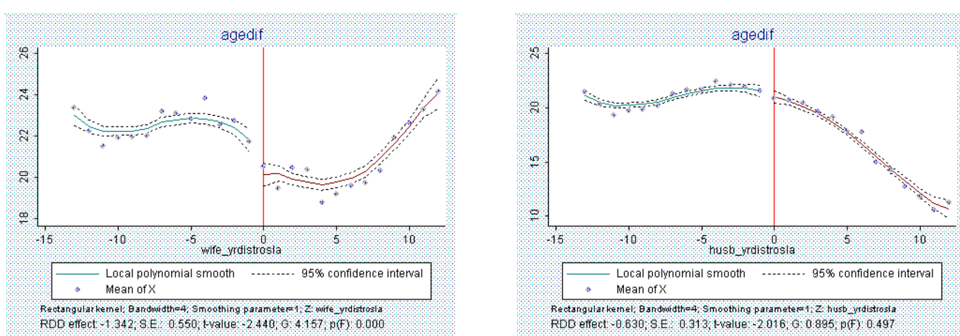

(a) Wives–Census (b) Husbands–Census

Notes: The graphs show local polynomial smooths of marital age gaps in months (husband-wife). The dots reflect mean age gaps by each bin on the abscissa. The solid line is the local polynomial smooth with the bandwidth and degree as shown using a rectangular kernel and the grid on the abscissa. The dashed lines are 95% confidence intervals of the local polynomial. 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, in particular Chris Marshall, Rachel Stuchbury and Wei Xun, as well as User Support (CeLSIUS). CeLSIUS is supported by the ESRC Census of Population Programme (Award Ref: RES-348-25-0004). The authors alone are responsible for the interpretation of the data. Census output is Crown copyright and is reproduced with the permission of the Controller of HMSO and the Queen's Printer for Scotland. Source: Census.

Table C.1: RoSLA on own academic qualification by month of birth

	(1)	(2)	(3)	(4)	(5)	(6)
	Married Women					
Optimal Bandwidth (24)	0.121***	0.114***	0.055***	0.120***	0.113***	0.054***
	(0.016)	(0.020)	(0.015)	(0.016)	(0.020)	(0.015)
N	40,264	40,264	40,264	40,264	40,264	40,264
G-statistic (p-value)	0.000	0.000	0.000	0.000	0.000	0.000
AIC	51,247	51,248	51,200	51,221	51,222	51,174
1/2 x Optimal B. (12)	0.089***	0.052**	0.096***	0.089***	0.052**	0.095***
	(0.013)	(0.015)	(0.017)	(0.013)	(0.015)	(0.017)
N	20,119	20,119	20,119	20,119	20,119	20,119
G-statistic (p-value)	0.004	0.086	0.277	0.002	0.044	0.165
AIC	25,689	25,679	25,675	25,675	25,665	25,661
2 x Optimal B. (48)	0.133***	0.124***	0.111***	0.133***	0.123***	0.110***
	(0.013)	(0.018)	(0.017)	(0.013)	(0.018)	(0.017)
N	79,397	79,397	79,397	79,397	79,397	79,397
G-statistic (p-value)	0.000	0.000	0.000	0.000	0.000	0.000
AIC	100,418	100,414	100,400	100,309	100,305	100,290
	Married Men					
Optimal Bandwidth (24)	0.111***	0.102***	0.092***	0.111***	0.101***	0.092***
	(0.013)	(0.020)	(0.025)	(0.013)	(0.020)	(0.025)
N	38,383	38,383	38,383	38,383	38,383	38,383
G-statistic (p-value)	0.000	0.000	0.000	0.000	0.000	0.000
AIC	51,625	51,628	51,622	51,590	51,593	51,587
1/2 x Optimal B. (12)	0.086***	0.112***	0.163***	0.086***	0.111***	0.163***
	(0.019)	(0.020)	(0.016)	(0.019)	(0.020)	(0.016)
N	19,391	19,391	19,391	19,391	19,391	19,391
G-statistic (p-value)	0.002	0.072	0.723	0.001	0.036	0.586
AIC	26,217	26,205	26,194	26,203	26,191	26,180
2 x Optimal B. (48)	0.111***	0.121***	0.106***	0.111***	0.121***	0.105***
	(0.011)	(0.014)	(0.017)	(0.011)	(0.014)	(0.017)
N	74,843	74,843	74,843	74,843	74,843	74,843
G-statistic (p-value)	0.000	0.000	0.000	0.000	0.000	0.000
AIC	100,029	100,021	100,020	99,929	99,922	99,921
Polyn. degree	1	2	3	1	2	3
Basic controls	No	No	No	Yes	Yes	Yes

Notes: The table shows estimates from local parametric estimation of equation (4.1) as described in Section 4.3.1 using different bandwidths, over rows, and polynomial degrees 1 to 3 over columns. The dependent variable is the spousal age gap measured in months (male-female). The bandwidth reflects the number of values of the running variable (month of birth) on each side of the discontinuity. Standard errors are robust and allow for random and identical specification errors. Below the estimates the p-value of the Lee and Card (2008) G-statistic and the Akaike Information Criterion indicate goodness of the polynomial fit. Robust standard errors reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table C.2: Change in marital qualifications gap at RoSLA threshold

	(1)	(2)	(3)	(4)	(5)	(6)
	Married Women					
Optimal Bandwidth (24)	-0.117***	-0.123***	-0.065*	-0.117***	-0.123***	-0.065*
	(0.016)	(0.019)	(0.027)	(0.016)	(0.019)	(0.027)
N	39,764	39,764	39,764	39,764	39,764	39,764
p_G	0.000	0.000	0.001	0.000	0.000	0.000
AIC	63,544	63,547	63,522	63,530	63,534	63,509
1/2 x Optimal B. (12)	-0.100***	-0.056	-0.141***	-0.100***	-0.057	-0.142***
	(0.017)	(0.031)	(0.019)	(0.017)	(0.031)	(0.018)
N	19,885	19,885	19,885	19,885	19,885	19,885
p_G	0.000	0.006	0.183	0.000	0.002	0.115
AIC	31,787	31,779	31,767	31,783	31,776	31,764
2 x Optimal B. (48)	-0.122***	-0.123***	-0.110***	-0.122***	-0.123***	-0.110***
	(0.012)	(0.017)	(0.017)	(0.012)	(0.017)	(0.017)
N	78,318	78,318	78,318	78,318	78,318	78,318
p_G	0.000	0.000	0.000	0.000	0.000	0.000
AIC	124,159	124,162	124,161	124,159	124,161	124,160
	Married Men					
Optimal Bandwidth (24)	0.108***	0.117***	0.090*	0.108***	0.117***	0.090*
	(0.016)	(0.025)	(0.036)	(0.016)	(0.025)	(0.036)
N	38,041	38,041	38,041	38,041	38,041	38,041
p_G	0.000	0.000	0.000	0.000	0.000	0.000
AIC	59,364	59,368	59,357	59,364	59,367	59,356
1/2 x Optimal B. (12)	0.094***	0.118***	0.162***	0.094***	0.118***	0.162***
	(0.023)	(0.030)	(0.034)	(0.023)	(0.030)	(0.034)
N	19,217	19,217	19,217	19,217	19,217	19,217
p_G	0.007	0.043	0.155	0.003	0.021	0.085
AIC	30,158	30,153	30,149	30,157	30,151	30,148
2 x Optimal B. (48)	0.095***	0.118***	0.114***	0.095***	0.118***	0.114***
	(0.012)	(0.017)	(0.022)	(0.012)	(0.017)	(0.022)
N	74,150	74,150	74,150	74,150	74,150	74,150
p_G	0.000	0.000	0.000	0.000	0.000	0.000
AIC	115,507	115,500	115,500	115,510	115,504	115,504
Polyn. degree	1	2	3	1	2	3
Basic controls	No	No	No	Yes	Yes	Yes

Notes: The table shows estimates from local parametric estimation of equation (4.1) as described in Section 4.3.1 using different bandwidths, over rows, and polynomial degrees 1 to 3 over columns. The dependent variable is the spousal age gap measured in months (male-female). The bandwidth reflects the number of values of the running variable (month of birth) on each side of the discontinuity. Standard errors are robust and allow for random and identical specification errors. Below the estimates the p-value of the Lee and Card (2008) G-statistic and the Akaike Information Criterion indicate goodness of the polynomial fit. Robust standard errors reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table C.3: Sensitivity to upper age limit

	No upper limit			Upper age=50			Upper age=45			Upper age=40		
Polynomial order	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
	Married Women											
Optimal Bandwidth (24)	-3.059***	-2.534**	-5.835***	-3.054***	-2.537**	-5.836***	-3.223***	-2.791**	-6.280***	-3.821***	-3.530**	-7.446***
	(0.696)	(0.825)	(1.326)	(0.695)	(0.825)	(1.326)	(0.763)	(0.876)	(1.404)	(0.833)	(1.141)	(1.359)
N	48668	48668	48668	48657	48657	48657	44525	44525	44525	37420	37420	37420
G-statistic (p-value)	0.022	0.016	0.088	0.022	0.016	0.089	0.006	0.004	0.029	0.016	0.010	0.084
1/2 x Optimal B. (12)	-4.595***	-3.898**	-0.674	-4.595***	-3.898**	-0.674	-4.889***	-4.417**	-0.347	-5.448***	-5.934***	-2.667
	(0.790)	(1.057)	(1.264)	(0.790)	(1.057)	(1.264)	(0.865)	(1.226)	(1.626)	(1.033)	(1.256)	(1.997)
N	24273	24273	24273	24273	24273	24273	22237	22237	22237	18686	18686	18686
G-statistic (p-value)	0.493	0.451	0.619	0.493	0.451	0.619	0.250	0.231	0.449	0.341	0.266	0.329
2 x Optimal B. (48)	-2.434***	-2.972***	-3.289***	-2.410***	-2.989***	-3.238***	-2.504***	-2.964***	-3.663***	-2.876***	-3.590***	-4.404***
	(0.542)	(0.702)	(0.760)	(0.548)	(0.711)	(0.775)	(0.599)	(0.793)	(0.867)	(0.623)	(0.876)	(1.045)
N	97424	97424	97424	96657	96657	96657	89047	89047	89047	74957	74957	74957
G-statistic (p-value)	0.025	0.023	0.020	0.013	0.012	0.011	0.002	0.002	0.001	0.004	0.004	0.004
	Married Men											
Optimal Bandwidth (24)	1.171	1.159	2.040*	1.170	1.160	2.040*	1.310*	1.206	2.186*	1.178	1.163	1.593
	(0.594)	(0.698)	(0.856)	(0.594)	(0.697)	(0.855)	(0.649)	(0.719)	(0.949)	(0.779)	(0.924)	(1.205)
N	44475	44475	44475	44466	44466	44466	40497	40497	40497	33452	33452	33452
G-statistic (p-value)	0.823	0.871	0.850	0.823	0.871	0.850	0.507	0.663	0.626	0.138	0.293	0.267
1/2 x Optimal B. (12)	2.087**	1.066	-1.302	2.087**	1.066	-1.302	2.068**	1.164	-0.773	1.829*	0.572	-1.566
	(0.691)	(0.817)	(1.158)	(0.691)	(0.817)	(1.158)	(0.646)	(0.775)	(0.885)	(0.863)	(1.018)	(1.082)
N	22373	22373	22373	22373	22373	22373	20393	20393	20393	16886	16886	16886
G-statistic (p-value)	0.814	0.826	0.872	0.814	0.826	0.872	0.861	0.919	0.947	0.581	0.736	0.804
2 x Optimal B. (48)	1.703***	1.116	1.430*	1.668***	1.083	1.485*	1.926***	1.109	1.518*	2.014**	0.951	1.542
	(0.444)	(0.647)	(0.706)	(0.443)	(0.646)	(0.712)	(0.469)	(0.670)	(0.715)	(0.597)	(0.843)	(0.978)
N	87928	87928	87928	87187	87187	87187	79870	79870	79870	65894	65894	65894
G-statistic (p-value)	0.475	0.448	0.475	0.456	0.430	0.449	0.225	0.223	0.275	0.010	0.012	0.023

Notes: The table shows estimates from local parametric estimation of equation (4.1) as described in Section 4.3.1 using different bandwidths, over rows, polynomial degrees 1 to 3 over columns. The dependent variable is the spousal age gap measured in months (male-female). The tables compare samples constructed using differing upper age limits as indicated. The bandwidth reflects the number of values of the running variable (month of birth) on each side of the discontinuity. Standard errors are robust and allow for random and identical specification errors. Below the estimates the p-value of the Lee and Card (2008) G-statistic indicate goodness of the polynomial fit. Robust standard errors reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

C.2 Technical Annex

C.2.1 The Regression Discontinuity Design

Eligibility for the increased compulsory school-leaving age was determined by a single date, whether an individual was born prior to or after 1st September 1957. Therefore the RoSLA reform can be thought of as a natural experiment inducing exogenous variation in the length of education received by an individual. As this variation was determined solely by a discontinuous function of an observed covariate, the individual's birth date, the estimation proceeds through a regression discontinuity design (RDD), an approach which allows the identification of causal treatment effects in quasi-experimental settings. The RDD approach is based on the idea that individuals are deterministically assigned to a treatment based on whether the value of an observed covariate, the running variable, Z_i, falls on either side of a threshold value $Z_i = z^*$ such that a discontinuity in the assignment function is induced precisely at this threshold value. The intuition is that individuals in the neighborhood of this threshold are identical in all other characteristics, apart from whether or not they are assigned to the treatment. Therefore by comparing individuals 'close' to the discontinuity from either side of the threshold, a causal effect of the treatment can be identified. As there is local randomisation additional covariates are not necessary, but may improve precision of the estimates (Lee and Lemieux, 2010).

C.2.2 Non-parametric Estimation

The estimate of interest is $\alpha^{RDD} = [(Y_1 - Y_0)|Z = z^*]$, the difference between the population means of individuals subject to the reform, Y_1, and individuals born prior to the reform implementation threshold, Y_0 evaluated at the threshold $Z = z^*$. With a continuously distributed running variable α^{RDD} is estimated by taking the difference in the limits of expected values in the outcome variable either side of the threshold so that

$$\hat{\alpha}^{RDD} = \lim_{e \to 0} E[Y_1|X = x^*] - E[Y_0|X = x^* - e] \approx E[Y_1 - Y_0|X = x^*] \tag{A1}$$

with non-parametric methods appropriate (Hahn et al., 2001) to be used in the estimation .

Visual depictions of the RDD analyses are produced via kernel-weighted local polynomial smoothing of the outcome variable either side of the RoSLA threshold. Although Fan and Gijbels (1996) advocate the use of triangular kernels due to their superior boundary properties, the estimation is not practically affected by kernel choice, and therefore a rectangular kernel is used as this has become the *de facto* standard in the literature (Imbens and Lemieux, 2008).

C.2.3 Parametric Estimation

Although the underlying distribution of the running variable (timing of birth) in the analysis is continuous, the Labour Force Survey contains only discrete measures of birth dates - month and year of birth. Therefore it is only possible to observe $E[Y_1|X = x^*]$, the outcome of an individual precisely at the threshold who is subject to the higher compulsory school-leaving age, and $E[Y_0|X = x^* - e]$, the outcome of individuals born before the threshold date. Furthermore, as the Z is discretely measured e is also discrete and hence takes on a finite number of values in $Z = z_j$ with $(j = 1, ..., J)$. Thus the closest realisation in the data below the threshold is $E[Y_0|Z = z^* - e] = E[Y_0|Z = z^* - 1]$, as it is not possible to allow $\lim_{e \to 0}$ as in (A1). Instead, in order to project the limit points in order to predict $E[Y_0|Z = z^* - 0]$ a parametric approach is required, and the estimation follows the methodology proposed by Lee and Card (2008) as described in Section C.2.1. Key choices affecting the validity of the estimates are the degree of the polynomial function applied and the size of the window width of observations around the discontinuity.

C.2.3.1 Polynomial Choice

In RDDs the magnitude of the estimates can be sensitive to the choice of polynomial in the running variable applied. Although a certain degree of smoothing minimises the influence of outliers and seasonality, an inappropriate order of polynomial in the empirical model may prove a poor approximation to the underlying data generating process. Information criteria such as the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) are often used in the literature to determine the optimal polynomial fit. Lee and Card (2008) propose a simple test statistic for choosing the best approximation of the data generating process, based on the closeness of the estimated polynomial function to the true distribution of the running variable Z. With a discrete running variable there is a natural comparison based on the smallest possible local polynomial, which captures the means calculated at each value of z_j. An auxiliary regression using a dummy for each value of z_j, which perfectly resembles means at each z_j is used to estimate this 'best' possible fit, which can then be compared to the fit of the polynomial applied in the estimation using the Lee and Card (2008) G-statistic:

$$G \equiv \frac{(RSS_R - RSS_{UR})/(J - K)}{RSS_{UR}/(N - J)}$$

where RSS_R is the residual sum of squares for the estimated model, and RSS_{UR} the residual sum of squares for the unrestricted model (the auxiliary regression). Assuming normality, G follows an $F_{(J-K,N-J)}$ distribution, with K the number of parameters estimated in the restricted model, N the number of observations and J the total number of individual values in the support of Z. The null hypothesis is that there is no systematic difference in the residual sum of squares in the restricted and unrestricted estimations, in other words that the order of polynomial used provides a good approximation of the underlying data generating process.

C.2.3.2 Window Width

A fundamental issue in RDD estimation is the size of the window of observations to use in the analysis. Theoretically the RD logic only needs a very small window of observations, as the treatment effect is produced through a comparison of the expectations just before and just after the discontinuity. However the smaller the window, the fewer the number of data points used in the estimation, which may lead to inefficient, although unbiased estimates. Increasing the window size increases the number of observations used in the analysis, thereby increasing the efficiency of the estimate. However too large a window may indicate that there is a discontinuity even if the data is in actuality smoothly distributed around the threshold. Therefore when choosing the appropriate window of observations an inherent trade-off arises between the bias and the efficiency of the estimate. Window width choice is thus equivalent to bandwidth choice in a non-parametric RDD with local linear regressions, and the parametric regressions therefore resemble local linear regressions of equal bandwidth around the threshold value.

C.2.3.3 Cross-Validation

Figure C.3: Cross-Validation Function

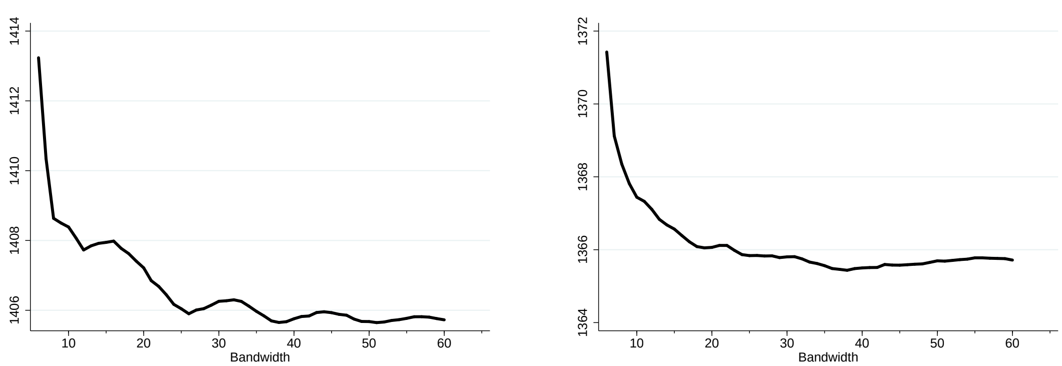

(a) Wives

(b) Husbands

Notes: The graphs show the cross validation function for the marital age gap, calculated for each spouse over bandwidths between 6 and 60.

The analysis uses the procedure proposed by Ludwig and Miller (2007) for optimal bandwidth selection in RDD settings. The cross-validation function is computed via a leave-one-out procedure. For each candidate bandwidth h in turn a set of regressions are computed: for every data point, i, a regression is estimated which omits observation i and a predicted value of the outcome variable is obtained. As RDD estimates are estimated at the boundary, where the value of the running variable for data point i is to the left of the threshold only observations where $z_i - h \leq z < z_i$ are used in the regression, and similarly if i lies to the right of the threshold then the regression uses only observations where $z_i < z \leq z_i + h$. The difference is then calculated between the predicted value from the regression which omits i, $\hat{y}(z_i)$, and the actual value y_i, which forms a value of the cross-validation procedure for bandwidth h. By repeating this procedure for each observation i over every possible bandwidth h and aggregating yields the cross-validation function $CV_Y(h) = \frac{1}{N}\sum_{i=1} N(y_i - \hat{y}(z_i))^2$. The optimal bandwidth is then determined by the value of h that minimises $CV_Y(h)$, the mean square difference of the predicted value to the true value of Y (Imbens and Lemieux (2008)).

The cross-validation function computed to determine the optimal bandwidth used in the estimation of the impact of the RoSLA reform on the marital age-gap is displayed in Figure C.3. For both husbands and wives the optimal bandwidth is indicated at approximately 24 months. Although the value of the cross-validation function continues to decline with as the bandwidth size increases, the incremental decrease in precision is marginal in comparison to the potential degree of bias in the estimate which may ensue from estimation with higher bandwidths.

Figure C.4: Sensitivity of Estimates to Bandwidth Choice

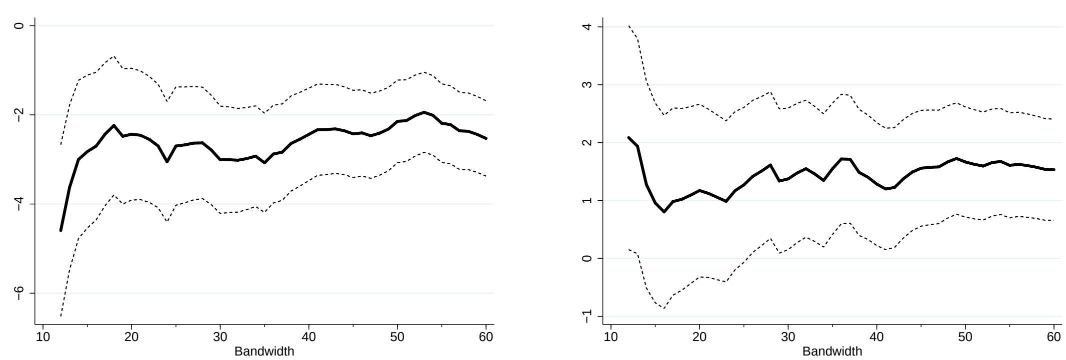

(a) Wives (b) Husbands

Notes: The graphs show RD estimates produced for each spouse of the effect of the RoSLA on the marital age gap as described in Section C.2.1 for each bandwidth choice over the range 12-60, using a quadratic polynomial in the running variable. The solid line indicates the magnitude of the estimate, the dashed lines indicate the 95% confidence interval.

C.2.4 Sensitivity of the Estimates

Direct examination of the sensitivity of the estimates to bandwidth choice is analogous to the cross-validation procedure described in Section C.2.3.3. Figure C.4 displays the magnitude of the estimates of the impact of the RoSLA reform on the marital age-gap estimated using bandwidths ranging from 12 to 60

As would be expected, both graphs indicate imprecise estimates with very low bandwidths, with the width of the confidence interval decreasing as bandwidth size increases. From a bandwidth size of approximately 24 months the magnitude of the coefficients becomes stable, and the increase in precision as indicated by the width of the confidence interval becomes small relative to the increase in the size of the bandwidth.

D Appendix to Chapter 5

D.1 Proofs

Proof of Lemma 1. We start by noting that, due to the functional form, $M(\pi_h, \pi_w, \epsilon, \hat{\phi})$ is a continuously differentiable function of $(\pi_h, \pi_w, \hat{\phi})$ and $D(\pi_w)$ is a continuously differentiable function of π_w. Differentiating yields that $\partial M/\partial\pi_h < 0$, $\partial D/\partial\pi_h = 0$, $\partial M/\partial\pi_w < 0$, and $\partial D/\partial\pi_w < 0$, and, importantly,

$$\frac{\partial\left(M-D\right)}{\partial\pi_h} < 0 \quad \text{and} \quad \frac{\partial\left(M-D\right)}{\partial\pi_w} > 0, \tag{a1}$$

where the latter inequality follows from concavity of $v\left(\cdot\right)$. Hence an increase in the wife's unemployment risk makes marriage more attractive to her, as the loss in earnings associated with unemployment has a larger negative impact on her utility when she does not have access to her partner's income.

Next we define

$$\pi_h' \equiv \begin{cases} 0 & \text{if } M\left(0,0,0,1\right) \leq D\left(0\right) \\ \sup\left\{\pi_h \in [0,1] \left| M\left(\pi_h, 0, 0, 1\right) \geq D\left(0\right)\right.\right\} & \text{if } M\left(0,0,0,1\right) > D\left(0\right) \end{cases} \tag{a2}$$

and

$$\pi_h'' \equiv \begin{cases} 1 & \text{if } M\left(1,1,0,1\right) \geq D\left(1\right) \\ \inf\left\{\pi_h \in [0,1] \left| M\left(\pi_h, 1, 0, 1\right) \leq D\left(1\right)\right.\right\} & \text{if } M\left(1,1,0,1\right) < D\left(1\right) \end{cases} \tag{a3}$$

Consider the case where $M\left(0,0,0,1\right) > D\left(0\right)$, the second case in (a2). By assumption A1, $M\left(1,0,0,1\right) < D\left(0\right)$. Hence it follows that $\pi_h' \in (0,1)$ and is the unique critical value for π_h at which $M = D$ given $\pi_w = 0$ (and $\epsilon = 0$ and $\hat{\phi} = 1$). Similarly, consider the case where $M\left(1,1,0,1\right) < D\left(1\right)$, the second case in (a3). By assumption A2, $M\left(0,1,0,1\right) > D\left(1\right)$. Hence it follows that $\pi_h'' \in (0,1)$ and is the unique critical value for π_h at which $M = D$ given $\pi_w = 1$ (and $\epsilon = 0$ and $\hat{\phi} = 1$). Next we verify that $\pi_h' < \pi_h''$. This follows trivially if $\pi_h' = 0$ and/or $\pi_h'' = 1$. Hence consider the case where $\pi_h' > 0$ and $\pi_h'' < 1$ (as in Figure 5.1). Note that since, per definition of π_h', $M\left(\pi_h', 0, 0, 1\right) = D\left(0\right)$, and using (a1) it follows that $M\left(\pi_h', 1, 0, 1\right) > D\left(1\right)$ and hence that $\pi_h'' > \pi_h'$.

Next we verify that (5.7) has a solution in the unit interval if and only if $\pi_h \in [\pi_h', \pi_h'']$. Consider the case where $\pi_h' > 0$. Then, $M\left(\pi_h, \pi_w, 0, 1\right) > D\left(\pi_w\right)$ at any $\left(\pi_h, \pi_w\right) \in [0, \pi_h') \times [0,1]$, implying that (5.7) does not have a solution in the unit interval. Similarly, consider the case where $\pi_h'' < 1$. Then, $M\left(\pi_h, \pi_w, 0, 1\right) < D\left(\pi_w\right)$ for any $\left(\pi_h, \pi_w\right) \in (\pi_h'', 1] \times [0,1]$, implying that (5.7) does not have a solution in the unit interval. Thus

(5.7) can have a solution in the unit interval only if $\pi_h \in [\pi_h', \pi_h'']$. Consider then some $\pi_h \in (\pi_h', \pi_h'')$. By definition of π_h' and π_h'' if follows that $M(\pi_h, 0, 0, 1) < D(0)$ and $M(\pi_h, 1, 0, 1) > D(1)$. It then follows from continuity of the value functions and (a1) that (5.7) has a unique solution we denote by $\hat{\pi}_w(\pi_h) \in (0, 1)$.

Implicitly differentiating (5.7) yields that

$$\frac{\partial \hat{\pi}_w}{\partial \pi_h} = -\frac{\partial (M - D) / \partial \pi_h}{\partial (M - D) / \partial \pi_w} > 0, \tag{a4}$$

where the sign follows from (a1).

The sign of the derivatives of $\hat{\pi}_w(\pi_h)$ with respect to the partners' wages follow in a similar way from the observation that

$$\frac{\partial (M - D)}{\partial \omega_h} > 0 \quad \text{and} \quad \frac{\partial (M - D)}{\partial \omega_w} < 0, \tag{a5}$$

where the latter inequality follows due to concavity of $v(\cdot)$. $\square$

Proof of Proposition 1. We first define the husband's expected utility in the case of divorce,

$$D(\pi_h, \epsilon) \equiv \mathbf{E}\left[u_h^d | \pi_h\right] - \alpha_h - \xi\epsilon, \tag{a6}$$

where $\mathbf{E}\left[u_h^d | \pi_h\right]$ is defined analogously to (5.4). The husband's expected utility from continued marriage on the other hand is type-dependent,

$$M(\pi_h, \pi_w, \epsilon; \theta) = \mathbf{E}\left[u_h^m | (\pi_h, \pi_w)\right] - \delta_\theta \kappa(\theta, \epsilon) - \xi\epsilon, \tag{a7}$$

where $\mathbf{E}\left[u_h^m | (\pi_h, \pi_w)\right]$ is defined analogously to (5.6). In particular, we obtain that a husband of type N ranks the possible outcomes with respect to marriage and behavioural effort in the following way:

$$M(\pi_h, \pi_w, 1; N) > M(\pi_h, \pi_w, 0; N) > D(\pi_h, 0) > D(\pi_h, 1). \tag{a8}$$

To see this, note that the first inequality follows from part (i) of assumption A4, the second inequality follows from part (ii) of assumption A4, and the third inequality is trivial. In contrast, a husband of type V ranks the possible outcomes in the following way:

$$M(\pi_h, \pi_w, 0; V) > M(\pi_h, \pi_w, 1; V) > D(\pi_h, 0) > D(\pi_h, 1). \tag{a9}$$

The first inequality follows from the assumption that $\delta_V = 0$. The second inequality follows from the fact that $\alpha_h > \xi$ which is implied by the combination of parts (i) and (ii) of assumption A4.

The key difference between (a8) and (a9) is that a husband of type V does not value the reduction in the risk of violence associated with the effort $\epsilon = 1$ whereas a husband of type N values it more than its cost.

There are four possible pure strategy profiles that the husband can adopt:

- **Strategy profile (1):** separation with $(\epsilon', \epsilon'') = (0, 1)$;
- **Strategy profile (2):** separation with $(\epsilon', \epsilon'') = (1, 0)$;
- **Strategy profile (3):** pooling with $(\epsilon', \epsilon'') = (1, 1)$;
- **Strategy profile (4):** pooling with $(\epsilon', \epsilon'') = (0, 0)$.

We will consider each possible pure strategy profile within each regime.

D.1.1 Regime R_1

Given that $(\pi_h, \pi_w) \in R_1$, the wife obtains a higher expected payoff from marriage than from divorce with any husband of type θ and any effort choice ϵ by the husband. We will now consider the four possible pure strategy profiles in turn:

Strategy profile (1). Bayesian updating implies that $\hat{\phi}(0) = 1$ and $\hat{\phi}(1) = 0$, and the wife rationally chooses to remain married at either choice of ϵ, $\chi' = \chi'' = m$. According to (a8) and (a9) each type of husband obtains his most preferred outcome and hence has no incentive to deviate, confirming that this is a PBE.

Strategy profile (2). Bayesian updating implies that $\hat{\phi}(0) = 0$ and $\hat{\phi}(1) = 1$, and the wife rationally chooses to remain married at either choice of ϵ, $\chi' = \chi'' = m$. In this case neither type of husband obtains his most preferred outcome and, since the wife responds to either choice of ϵ by continuing the marriage, each type of husband would have an incentive to deviate.

Strategy profile (3). Bayesian updating implies that $\hat{\phi}(1) = \phi$, while $\hat{\phi}(0)$ is not determined by Bayesian updating. Irrespective of how the wife updates her beliefs at $\epsilon = 0$, she rationally chooses to remain married at either choice of ϵ, $\chi' = \chi'' = m$. Given this, a husband of type V would be better off deviating to $\epsilon = 0$.

Strategy profile (4). Bayesian updating implies that $\hat{\phi}(0) = \phi$, while $\hat{\phi}(1)$ is not determined by Bayesian updating. Irrespective of how the wife updates her beliefs at $\epsilon = 1$, she rationally chooses to remain married at either choice of ϵ, $\chi' = \chi'' = m$. Given this, a husband of type N would be better of deviating to $\epsilon = 1$.

D.1.2 Regime R_0

In this regime, the wife's decision whether or not to remain married depends on her beliefs and on the husband's observed effort.

Strategy profile (1). Bayesian updating implies that $\hat{\phi}(0) = 1$ and $\hat{\phi}(1) = 0$. The wife then (by assumptions A1 and A3) continues the marriage if and only if the husband makes the effort $\epsilon = 1$, that is $\chi'' = m$ and $\chi' = d$. A type V would then be better of deviating to $\epsilon = 1$ as by doing so he would avoid triggering divorce.

Strategy profile (2). Bayesian updating implies that $\hat{\phi}(0) = 0$ and $\hat{\phi}(1) = 1$. Given these updated beliefs, the wife rationally responds (by Assumption A3) to $\epsilon = 0$ by continuing the marriage, that is $\chi' = m$. This then cannot be an equilibrium since a type V husband could then deviate to $\epsilon = 0$ and obtain is his most preferred outcome.

Strategy profile (3). Bayesian updating implies that $\hat{\phi}\,(1) = \phi$ and, by assumption A3, the wife rationally responds to $\epsilon = 1$ by continuing the marriage, $\chi'' = m$. Note that $\hat{\phi}\,(0)$ is not determined by Bayesian updating. Suppose that the wife, at $\epsilon = 0$, believes that the husband is of type V, that is $\hat{\phi}\,(0) = 1$. She would then rationally respond to $\epsilon = 0$ by choosing divorce, $\chi' = d$. Given this, and given the preference orderings in (a8) and (a9), neither husband type has any incentive to deviate. Note also that the out-of-equilibrium belief $\hat{\phi}\,(0) = 1$ satisfies the Choo-Kreps "intuitive criterion". For a husband of type N, $\epsilon = 0$ is equilibrium dominated as this type, by choosing $\epsilon = 1$, obtains his most preferred outcome in equilibrium. In contrast, a husband of type V would benefit if the wife were to respond to $\epsilon = 0$ by continuing the marriage.

Strategy profile (4). Bayesian updating implies that $\hat{\phi}\,(0) = \phi$ but does not determine $\hat{\phi}\,(1)$. Given this, and by assumption A3, the wife rationally continues the marriage upon observing $\epsilon = 0$, that is $\chi' = m$. Next, note that by (a8) for a husband of type N in particular to prefer to choose $\epsilon = 0$ it must be that the wife responds to $\epsilon = 1$ by divorcing, that is $\chi'' = d$. Hence for this to be a PBE, $\hat{\phi}\,(1)$ must be such that the wife prefers divorce upon observing $\epsilon = 1$. In particular, from Assumption 3 it must be that $\hat{\phi}\,(1) > \phi$. Such a PBE however does not satisfy the "intuitive criterion". For a husband of type V, $\epsilon = 1$ is equilibrium dominated as this type, by choosing $\epsilon = 0$, obtains his most preferred outcome in equilibrium. In contrast, a husband of type N would benefit from deviating if the wife were to respond to $\epsilon = 1$ by continuing the marriage. Hence, by the "intuitive criterion" the wife's out-of-equilibrium beliefs must be $\hat{\phi}\,(1) = 0$, contradicting that she would choose$\chi'' = d$. □

D.2 Appendix B: Variable Descriptions

The following variables are used in Section 5.4.3 ("Extended Results"):

1. **Magistrate court timeliness:** This is a measure of the duration from first listing of an offence to completion, for defendants in indictable cases in magistrates courts, and hence captures the "efficiency" of the criminal justice system, post arrest. The data is released on an annual basis from the Ministry of Justice, and is at the Local Justice Area (LJA) geography which coincides with the PFAs we use in the analysis.

2. **Police force manpower:** This variable refers to overall police manpower per 10,000 capita at PFA level. It is comprised of the number of (full-time equivalent) police officers, police community support officers, and police staff. This data is released annually by the Home Office.

3. **Violent crime rate:** This is the number of recorded violent crimes per 10,000 capita at PFA level. The data is from the Home Office.

4. **Non-violent crime rate:** This is the number of recorded non-violent crimes per 10,000 capita at PFA level. The data is from the Home Office.

5. **Alcohol hospitalizations:** This is the number of alcohol hospitalizations per 10,000 capita at PFA level. This is from the Local Alcohol Profiles for England datasets, available from the North West Public Health Observatory data, which is part of Public Health England. Note that this data is not available for the 4 welsh PFAs. We aggregated the data up to PFA level from Local Authority level.

6. **Internal migration:** These are the number of in- and out-migrants as a percentage of the PFA population in each age/gender group. The statistics are compiled using the data series "Internal Migration by Local Authorities in England and Wales" which are released annually by the Office for National Statistics (ONS) to coincide with the mid-year population estimates. The data has received the "National Statistics" accreditation, and are understood to be the best official source of information on internal migration in England and Wales. The data is available by gender and in 5 year age groups at Local Authority level. Here we aggregated up to PFA level and using the APS defined age grouping.

7. **Drugs possession:** This is the number of arrests for possession per 10,000 capita at PFA level. This data is from the quarterly Home Office Offences tables.

The data in (1)-(6) come from annual tables, so has been interpolated to produce data at the period frequency.

D.3 Complete Set of Estimated Marginal Effects

Table D.1: Impact of Unemployment on Physical Abuse - Full Set of Results from Main Specification

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
Unemployment	-0.026								0.008
in own-age group	(0.018)								(0.019)
Female unemployment		0.097**	0.093**	0.102**	0.094**	0.083**	0.102**	0.091**	
in own-age group		(0.027)	(0.027)	(0.028)	(0.027)	(0.029)	(0.028)	(0.035)	
Male unemployment		-0.090**	-0.097**	-0.081**	-0.089**	-0.094**	-0.067*	-0.084*	
in own-age group		(0.021)	(0.022)	(0.027)	(0.021)	(0.029)	(0.029)	(0.037)	
Female unemployment			-0.013						
in other age groups			(0.065)						
Male unemployment			-0.047						
in other age groups			(0.055)						
Female real wage				0.005					
in own-age group				(0.009)					
Male real wage				-0.001					
in own-age group				(0.006)					
Unemployment rate gap									0.094**
(F-M) in own age group									(0.022)
Age in years	-0.001**	-0.001**	-0.001**	-0.001**	-0.001**				-0.001**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)				(0.000)

continued on next page

continued from previous page

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
Age Group: 16-24						0.039**	0.027	0.044**	
						(0.004)	(0.014)	(0.017)	
Age Group: 25-34						0.027**	0.013	0.009	
						(0.002)	(0.013)	(0.014)	
Age Group: 35-49						0.015**	0.009	0.002	
						(0.002)	(0.014)	(0.014)	
Ethnicity: White	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Ethnicity: Asian	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Ethnicity: Black	-0.010	-0.010	-0.010	-0.010	-0.010	-0.010	-0.010*	-0.010*	-0.010
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Qualifications: Other	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Qualifications: GCSE	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003
grades A-C	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Qualifications: A Level	-0.009**	-0.009**	-0.009**	-0.009**	-0.009**	-0.008**	-0.009**	-0.009**	-0.009**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Qualifications: Higher educ,	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**

continued on next page

continued from previous page

Specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
below degree	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Qualifications: Degree or above	-0.020**	-0.020**	-0.020**	-0.020**	-0.020**	-0.020**	-0.020**	-0.020**	-0.020**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Religion: Christian	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Religion: Muslim	-0.009	-0.009	-0.009	-0.009	-0.009	-0.009	-0.010	-0.009	-0.009
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Religion: Hindu	-0.015	-0.015	-0.015	-0.015	-0.015	-0.016	-0.016	-0.016	-0.015
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.008)	(0.008)	(0.009)
Religion: Sikh	-0.011	-0.011	-0.011	-0.011	-0.011	-0.012	-0.012	-0.012	-0.011
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Religion: Jewish	-0.037*	-0.037*	-0.037*	-0.037*	-0.037*	-0.037*	-0.037*	-0.037*	-0.037*
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Religion: Buddhist	0.008	0.008	0.008	0.008	0.008	0.007	0.007	0.007	0.008
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Religion: Other	0.009	0.009	0.009	0.009	0.009	0.008	0.009	0.009	0.009
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Number of children	0.005**	0.005**	0.005**	0.004**	0.005**	0.005**	0.005**	0.005**	0.005**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)

continued on next page

continued from previous page

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
HH contains kids aged under 5	0.005*	0.005*	0.005*	0.005*	0.005*	0.005*	0.005*	0.005*	0.005*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Area and time fixed effects	yes	yes	yes	yes	yes	yes	yes	yes	yes
Linear age-in-years control	yes	yes	yes	yes	yes	no	no	no	yes
Age group fixed effects	no	no	no	no	no	yes	yes	yes	no
Age group * Period FEs	no	no	no	no	no	no	yes	yes	no
Age group * Areas FEs	no	no	no	no	no	no	no	yes	no
Other demographic controls	yes	yes	yes	yes	yes	yes	yes	yes	yes
Area-specific linear trends	no	no	no	no	yes	no	no	no	no
Obs.					86,731				

Notes: See notes to Table 4. ** Significant at 1%. * Significant at 5%.

Table D.2: Impact of Unemployment on Non-Physical Abuse - Full Set of Results from Main Specification

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
Unemployment	-0.012								0.021
in own-age group	(0.023)								(0.024)
Female unemployment		0.102**	0.108**	0.110**	0.104**	0.078*	0.093*	0.087	
in own-age group		(0.037)	(0.038)	(0.038)	(0.037)	(0.040)	(0.042)	(0.048)	
Male unemployment		-0.081**	-0.074*	-0.061	-0.085**	-0.090*	-0.066	-0.077	
in own-age group		(0.030)	(0.032)	(0.037)	(0.031)	(0.039)	(0.039)	(0.048)	
Female unemployment			0.031						
in other age groups			(0.080)						
Male unemployment			0.035						
in other age groups			(0.068)						
Female real wage				-0.002					
in own-age group				(0.010)					
Male real wage				0.008					
in own-age group				(0.007)					
Unemployment rate gap									0.093**
(F-M) in own age group									(0.032)
Age in years	-0.001**	-0.001**	-0.001**	-0.001**	-0.001**				-0.001**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)				(0.000)

continued on next page

continued from previous page

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
Age Group: 16-24						0.046**	0.027	0.043	
						(0.005)	(0.021)	(0.024)	
Age Group: 25-34						0.029**	0.021	0.008	
						(0.003)	(0.017)	(0.018)	
Age Group: 35-49						0.020**	0.018	0.025	
						(0.002)	(0.017)	(0.017)	
Ethnicity: White	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.005
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Ethnicity: Asian	-0.011	-0.011	-0.011	-0.011	-0.011	-0.011	-0.011	-0.011	-0.011
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Ethnicity: Black	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Qualifications: Other	0.000	0.000	0.000	-0.000	-0.000	-0.000	-0.000	-0.000	0.000
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Qualifications: GCSE	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003
grades A-C	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Qualifications: A Level	-0.009**	-0.009**	-0.009**	-0.010**	-0.009**	-0.009**	-0.009**	-0.009**	-0.009**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Qualifications: Higher educ,	-0.008**	-0.008**	-0.008**	-0.009**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**

continued on next page

continued from previous page

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
below degree	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Qualifications: Degree or above	-0.023**	-0.024**	-0.024**	-0.024**	-0.023**	-0.023**	-0.023**	-0.023**	-0.024**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Religion: Christian	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.008**	-0.009**	-0.008**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Religion: Muslim	-0.012	-0.013	-0.012	-0.013	-0.013	-0.012	-0.012	-0.012	-0.013
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Religion: Hindu	0.002	0.002	0.002	0.002	0.002	0.001	0.002	0.001	0.002
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Religion: Sikh	0.017	0.017	0.017	0.017	0.017	0.017	0.017	0.016	0.017
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Religion: Jewish	-0.021	-0.021	-0.021	-0.021	-0.022	-0.022	-0.022	-0.021	-0.021
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Religion: Buddhist	0.004	0.004	0.004	0.004	0.004	0.004	0.003	0.003	0.004
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Religion: Other	0.006	0.006	0.006	0.006	0.006	0.006	0.006	0.006	0.006
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Number of children	0.007**	0.007**	0.007**	0.007**	0.007**	0.007**	0.007**	0.007**	0.007**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)

continued on next page

continued from previous page

Specification	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
HH contains kids aged under 5	0.004	0.004	0.004	0.004	0.004	0.005*	0.004	0.004	0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)
Area and time fixed effects	yes	yes	yes	yes	yes	yes	yes	yes	yes
Linear age-in-years control	yes	yes	yes	yes	yes	no	no	no	yes
Age group fixed effects	no	no	no	no	no	yes	yes	yes	no
Age group * Period FEs	no	no	no	no	no	no	yes	yes	no
Age group * Areas FEs	no	no	no	no	no	no	no	yes	no
Other demographic controls	yes	yes	yes	yes	yes	yes	yes	yes	yes
Area-specific linear trends	no	no	no	no	yes	no	no	no	no
Observations					86,731				

Notes: See notes to Table 4. ** Significant at 1%. * Significant at 5%.

D.4 A Simple Model of Household Bargaining Under Uncertainty

In this appendix, we present a bargaining model of domestic violence. The model extends the Nash bargaining approach presented by Aizer (2010) to allow for income uncertainty. In order to simplify the analysis we assume additively separable preferences. When incomes are uncertain, the couple has an incentive to bargain at the ex-ante stage, before their incomes are realized, and we assume that the outcome of their ex-ante negotiations is binding.

As one would expect, a key feature of ex-ante bargaining is risk sharing. Hence the couple's ex-ante bargained allocation will smooth consumption as far as possible given the uncertainty they face regarding total household income. However, by direct analogy, the couple also have an incentive to "smooth violence" across states of nature. As there is no uncertainty regarding the available choices of violence, the ex-ante bargained allocation features equilibrium violence that is independent of the income realization. However, it is not independent of the partners' income *prospects*. Generalizing the theoretical prediction from Aizer (2010), we show that a shifting of the income probability distribution which reduces the husband's expected income and increases the wife's expected income while leaving the probability distribution over household income unchanged reduces the ex-ante bargained level of violence.

This conclusion holds for two possible consequences of failing to agree in the ex-ante bargaining. It holds if a failure to agree ex-ante implies that the couple will not engage in any further negotiations but instead behave non-cooperatively or divorce, and it also holds if failure to agree ex-ante leads to ex-post bargaining once all uncertainty is resolved.

D.4.1 Setup

Consider a couple consisting of a husband h and a wife w. Let the preferences of the spouses be defined over private consumption (c_i) and violence (v), with the husband's utility increasing in violence and the wife's decreasing in violence. For simplicity, suppose that the utility functions of the spouses are additively separable and given by

$$U_h(c_h, v) = u_h\left(c_h\right) + \varphi_h\left(v\right) \quad \text{and} \quad U_w(c_w, v) = u_w\left(c_w\right) + \varphi_w\left(1 - v\right), \tag{d1}$$

where $c_i \in R_+$ and $v \in [0,1]$, and where each sub-utility function is twice continuously differentiable, strictly increasing and strictly concave, with $u_i\left(c_i\right) \to -\infty$ as $c_i \to 0^+$.

Each partner faces income uncertainty, with y_h and y_w being independent draws from two distributions $F_h\left(y_h\right)$ and $F_w\left(y_w\right)$ defined on a common discrete support denoted $Y \equiv \{y_1, y_2, ..., y_N\}$, ordered increasingly. The associated probability density functions are denoted by $f_h(y_h)$ and $f_w(y_w)$, respectively. Hence the set of possible *states of the world* is $Y \times Y = Y^2$ with a typical element (y_h, y_w). The probability distributions are known to the couple who bargain ex-ante, before uncertainty is resolved, over which allocation to choose. An *allocation* is defined as a mapping $\{c_h(y_h, y_w), c_w(y_h, y_w), v(y_h, y_w)\}$ detailing the couple's consumption profile and violence choice in each state of the world $(y_h, y_w) \in Y^2$. The consumption profile (c_h, c_w) chosen at the state (y_h, y_w) must satisfy being non-negative in both components and $c_h + c_w \leq y_h + y_w$.

D.4.2 Ex-Ante Bargaining: Consumption and Violence Smoothing

When bargaining ex-ante, the fallback is either to bargain ex-post or not to bargain at all. If the fallback is not to bargain at all, then each partner j will have a fallback expected utility which depends only on his or her own income distribution F_j. If the fallback is to bargain ex-post—i.e., once all uncertainty has been resolved—then each partner's fallback expected utility depends on both F_h and F_w. Both cases will be considered below. We will highlight here some properties of ex-ante bargaining which are *independent* of the nature of the fallback. Hence we adopt the general notation $U_i^0(F)$ for the fallback expected utility of partner i, where $F \equiv \{F_h, F_w\}$.

Given an equilibrium-negotiated allocation $\{c_h(y_h, y_w), c_w(y_h, y_w), v(y_h, y_w)\}$, the gain in expected utility to the husband is

$$\Delta_h = U_h^* - U_h^0(F) = \sum_{y_h \in Y} \sum_{y_w \in Y} f_h(y_h) f_w(y_w) \left[u_h(c_h(y_h, y_w)) + \varphi_h(v(y_h, y_w))\right] - U_h^0(F), \tag{d2}$$

while the corresponding gain in expected utility to the wife is

$$\Delta_w = U_w^* - U_w^0(F) = \sum_{y_h \in Y} \sum_{y_w \in Y} f_h(y_h) f_w(y_w) \left[u_w(c_w(y_h, y_w)) + \varphi_w(1 - v(y_h, y_w))\right] - U_w^0(F), \tag{d3}$$

where U_h^* and U_w^* are the equilibrium expected utilities of the husband and the wife respectively.

The ex-ante Nash bargained agreement maximizes $\Delta_h \Delta_w$. Consider first the first order conditions with respect to the partners' consumption levels in state (y_h, y_w). These reduce to:

$$\frac{u_h'(c_h(y_h, y_w))}{u_w'(c_w(y_h, y_w))} = \Delta_r, \tag{d4}$$

where

$$\Delta_r \equiv \frac{\Delta_h}{\Delta_w}, \tag{d5}$$

denotes the relative expected utility gain of the husband. Noting that the right hand side of (d4) is independent of the state of the world, it follows that the same is true of the left hand side. Hence, as the bargained outcome is ex-ante efficient it features *complete consumption insurance* in the standard sense that the ratio of the partners' marginal utilities of consumption is constant across states of the world (see e.g. Cochrane, 1991). It does *not* imply complete consumption smoothing in the sense that each partner has an consumption that is independent of the state of the world: this is since the couple face uncertainty regarding total household income, $y_h + y_w$, which per construction is not constant across states of the world.

Considering violence, the first order condition for the bargained level of violence $v(y_h, y_w)$ reduces to

$$\frac{\varphi_h'(v(y_h, y_w))}{\varphi_w'(1 - v(y_h, y_w))} = \Delta_r. \tag{d6}$$

Noting again that the right hand side is constant across states of the world, it follows that the same is true for the left hand side. In contrast to consumption, this implies that $v(y_h, y_w)$ *is* constant across states of the world. The analogy to consumption is clear: in both cases, concavity of each partner's utility function implies a benefit from smoothing. In the case of consumption, the possibility for smoothing is limited due to the uncertainty about total household income. There is no such uncertainty regarding the available choices of violence, and thus violence is perfectly smoothed across states of the world. Hence the following conclusion holds irrespective of the specification of the fallback utilities.

Lemma 2. *Ex-ante Nash bargaining by the couple leads to:*

(a) Complete consumption insurance: the partners' relative marginal utilities are constant across states of the worlds [see eq. (d4)];

(b) Complete violence smoothing: the chosen violence level is constant across states of the world [see eq. (d6)].

Moreover, as can be seen from (d4) and (d6), the bargained outcome is effectively summarized by Δ_r. Of particular interest to us is to note that:

Lemma 3. *The ex-ante bargained state-independent level of violence* $v^* = v(y_h, y_w)$ *is strictly decreasing in* Δ_r.

In general, the ex-ante bargained allocation "discriminates" against the partner whose expected utility gain from implementing it exceeds that of the other partner. Thus, as the relative expected utility gain of the husband (Δ_r) increases, he has to "compensate" his spouse by agreeing to a lower level of equilibrium violence.

In order to conduct comparative statics on the bargained outcome, it is useful to rephrase the bargaining problem as the general problem of choosing expected utilities U_h^* and U_w^* for the two partners in order to maximize

$$\left(U_h^* - U_h^0(F)\right)\left(U_w^* - U_w^0(F)\right), \tag{d7}$$

subject to (U_h^*, U_w^*) being in a feasible set. In order to define the feasible set of expected utilities we first formally define the set of feasible allocations.

Definition 1. *An allocation* $\{c_h(y_h, y_w), c_w(y_h, y_w), v(y_h, y_w)\}$ *is said to be feasible if for all states of the world* $(y_h, y_w) \in Y^2$ *and for each* $i \in \{h, w\}$*:* $c_i(y_h, y_w) \in [0, y_h + y_w]$, $c_h(y_h, y_w) + c_w(y_h, y_w) \leq y_h + y_w$, *and* $v(y_h, y_w) \in [0, 1]$.

We can now define a feasible expected utility profile

Definition 2. *The expected utility profile* (U_h, U_w) *is said to be feasible if there exists a feasible allocation* $\{c_h(y_h, y_w), c_w(y_h, y_w), v(y_h, y_w)\}$ *such that for each state of the world* $(y_h, y_w) \in Y^2$*:*

$$U_h = \sum_{y_h \in Y} \sum_{y_w \in Y} f_h(y_h) f_w(y_w) \left[u_h(c_h(y_h, y_w)) + \varphi_h(v(y_h, y_w))\right],$$

and

$$U_w = \sum_{y_h \in Y} \sum_{y_w \in Y} f_h(y_h) f_w(y_w) \left[u_w(c_w(y_h, y_w)) + \varphi_w(1 - v(y_h, y_w))\right].$$

The set of feasible expected utility profiles is denoted T. We want to demonstrate that T is a convex set. Let $\left(U_h^0, U_w^0\right)$ and $\left(U_h^1, U_w^1\right)$ be two elements in T. We then need to verify that, for any $\alpha \in (0, 1)$

$$\left(U_h^2, U_w^2\right) \equiv \left(\alpha U_h^0 + (1-\alpha) U_h^1, \alpha U_w^0 + (1-\alpha) U_w^1\right), \tag{d8}$$

is also in the set T. Let $\left\{c_h^k(y_h, y_w), c_w^k(y_h, y_w), v^k(y_h, y_w)\right\}$ denote a feasible allocation that supports the expected utility profile (U_h^k, U_w^k) for each $k = 0, 1$. Consider then the convex combination of the two supporting allocations: at each node (y_h, y_w) define

$$\hat{c}_i(y_h, y_w) = \alpha c_i^0(y_h, y_w) + (1-\alpha) c_i^1(y_h, y_w), \tag{d9}$$

for $i = h, w$, and

$$\hat{v}(y_h, y_w) = \alpha v^0(y_h, y_w) + (1-\alpha) v^1(y_h, y_w), \tag{d10}$$

and note that this is a feasible allocation. Consider then the expected utility profile generated by this allocation. For the husband we obtain the expected utility,

$$\hat{U}_h = \sum_{y_h \in Y} \sum_{y_w \in Y} f_h(y_h) f_w(y_w) \left[u_h(\hat{c}_h(y_h, y_w)) + \varphi_h(\hat{v}(y_h, y_w)) \right]. \tag{d11}$$

Due to concavity of $u_h(\cdot)$ and $\varphi_h(\cdot)$ it follows that, in each state of the world:

$$u_h(\hat{c}_h(y_h, y_w)) > \alpha u_h\left(c_i^0(y_h, y_w)\right) + (1-\alpha)\alpha u_h\left(c_i^1(y_h, y_w)\right), \tag{d12}$$

and

$$\varphi_h(\hat{v}(y_h, y_w)) > \alpha \varphi_h\left(v^0(y_h, y_w)\right) + (1-\alpha)\varphi_h\left(v^1(y_h, y_w)\right), \tag{d13}$$

and hence it follows that $\hat{U}_h > U_h^2$. An identical argument shows that, for the wife, $\hat{U}_w > U_w^2$. Since it is always possible to reduce the expected utility of either (or both partners) by reducing consumption at some arbitrary node, it follows that $\left(U_h^2, U_w^2\right) \in T$. Moreover, the argument above makes clear that if even if $\left(U_h^0, U_w^0\right)$ and $\left(U_h^1, U_w^1\right)$ are both boundary points of T, $\left(U_h^2, U_w^2\right)$ is not a boundary point. Hence we have that:

Lemma 4. *The feasible set of expected utilities T is strictly convex.*

We also take it as given that the set T is compact. For simplicity we further assume that the Pareto frontier—i.e., the downward sloping part of the boundary of T—is twice differentiable. Letting $U_w(U_h)$ denote the Pareto frontier, it thus follows that $U_w'(U_h) < 0$ and $U_w''(U_h) < 0$.

The solution to the ex ante bargaining problem (d7) satisfies the general first order condition

$$\Delta_r \equiv \frac{\left(U_h^* - U_h^0(F)\right)}{\left(U_w^* - U_w^0(F)\right)} = -\frac{1}{U_w'\left(U_h^*\right)}, \tag{d14}$$

where $U_w^* = U_w(U_h^*)$. This feature will be key to the comparative statics below.

D.4.3 Comparative Statics with Autarky ("Divorce") as the Threat Point

In order to conduct a comparative statics analysis, we specify the fallback to be autarky. Ex-post bargaining as a fallback (see e.g. Riddell, 1981) will be considered below. Hence we define the fallback utilities to be:

$$U_h^0(F_h) = \sum_{y_h \in Y} f_h(y_h)\left[u_h(y_h) + \varphi_h(0)\right] \quad \text{and} \quad U_w^0(F_w) = \sum_{y_w \in Y} f_w(y_w)\left[u_w(y_w) + \varphi_w(1)\right], \tag{d15}$$

for the husband and the wife respectively. Thus, when living in autarky each spouse consumes his or her own income and there is no violence.

Having assumed that the two partners have income distributions with the same support, we can now consider a simple comparative static exercise. Consider two income levels $\underline{y}$ and $\overline{y}$ in Y with $\overline{y} > \underline{y}$ and a small constant $\Delta > 0$. Then consider the following shifting of probability:

$$\Delta f_h\left(\underline{y}\right) = \Delta,\ \Delta f_h\left(\overline{y}\right) = -\Delta,\ \Delta f_w\left(\underline{y}\right) = -\Delta,\ \Delta f_w\left(\overline{y}\right) = \Delta. \tag{d16}$$

Hence there is a shifting of probability mass Δ for each partner. For the husband, this shifting involves decreasing the probability of the higher income level $\overline{y}$ and increasing the probability of the lower income level $\underline{y}$. For the wife, the shifting goes in the opposite direction.

In interpreting the model, we can think of the lower income level $\underline{y}$ as unemployment and the higher level $\overline{y}$ as employment. The perturbation thus increases the husband's probability of unemployment while increasing the wife's probability of employment. We will show that the shifting of probability leads to a reduction in the ex-ante bargained level of violence.

Note in particular that, per construction, the income shift in (d16) does not affect the distribution of household income. Hence the perturbation leaves the feasible set of expected utilities T unchanged.[1] Next we note that the perturbation decreases the fall-back/autarky value for the husband but increases it for the wife,

$$\Delta U_h^0\left(F_h\right) = \Delta\left[u_h\left(\underline{y}\right) - u_h\left(\overline{y}\right)\right] < 0 \text{ and } \Delta U_w^0\left(F_w\right) = -\Delta\left[u_w\left(\underline{y}\right) - u_w\left(\overline{y}\right)\right] > 0. \tag{d17}$$

Consider then the impact of the reform on the bargaining outcome, in particular on (d14). As the reform has not affected the set of feasible expected utility profiles, it has not changed the Pareto frontier $U_w\left(U_h\right)$. From inspecting (d14) we obtain the following key result:

Lemma 5. *The shifting of probability in eq. (d16) leads to:*

(a) A decrease in the husband's equilibrium expected utility U_h^**;*

(b) An increase in the wife's equilibrium expected utility U_w^**;*

(c) An increase in the relative expected utility gain of the husband $\Delta_r = \frac{U_h^* - U_h^0(F_h)}{U_w^* - U_w^0(F_w)}$.

[1] In principle, the argument for this requires the definition of a feasible allocation to be generalized to allow for randomization at any given state of the world. This means that if the couple behave differently at the two nodes $\left(\underline{y}, \overline{y}\right)$ and $\left(\overline{y}, \underline{y}\right)$, then after the shift in probability they can still "replicate" the same probability distribution over outcomes by adopting the behaviour associated with node $\left(\underline{y}, \overline{y}\right)$ at node $\left(\overline{y}, \underline{y}\right)$ with probability Δ.

The first two parts are intuitive results. The third part, which is central for our purposes, says that, as the husband's probability of unemployment increases, he has more to gain in expected utility terms than his spouse from striking an ex-ante agreement. As a consequence, his relative bargaining position weakens. Combining Lemmas (3) and (5) we obtain the main result:

Proposition 2. *Suppose that the relevant threat point in the ex-ante bargaining process is autarky ("divorce"). Then the shifting of probability in eq. (d16) leads to a decrease in the ex-ante bargained state-independent equilibrium level of violence* $v^* = v(y_h, y_w)$.

D.4.4 Comparative Statics with Ex-Post Bargaining as the Threat Point

The assumption of divorce in the case of failure to agree in ex-ante negotiations may be overly strong. If the couple cannot agree on an allocation at the ex-ante stage, they can still bargain ex-post once all uncertainty is resolved.[2] We show here that Proposition 2 also holds in this case. In order to demonstrate that result we need to start by characterizing the outcome of ex-post Nash bargaining over consumption levels and violence.

D.4.4.1 Ex-Post Bargaining

Suppose that the state of the world (y_h, y_w) has been realized without any ex-ante agreement having been reached. The couple can then bargain over the allocation of consumption ex post. The fallback position here is "no trade" (or divorce). Hence in absence of an agreement the partners' utilities are

$$U_h^0 = u_h(y_h) + \varphi_h(0) \text{ and } U_w^0 = u_w(y_w) + \varphi_w(1), \tag{d18}$$

respectively. Ex-post Nash bargaining solves $\max \Delta_h^* \Delta_w^*$ where

$$\Delta_h^* = U_h - U_h^0 = u_h(c_h) + \varphi_h(v) - U_h^0, \tag{d19}$$

and

$$\Delta_w^* = U_w - U_w^0 = u_w(c_w) + \varphi_w(v) - U_w^0,$$

and subject to feasibility, $c_h + c_w \leq y_h + y_w$ and $v \in [0,1]$. The first order conditions with respect to consumption and violence imply

$$\frac{u_h'(c_h)}{u_w'(c_w)} = \frac{\Delta_h}{\Delta_w}, \tag{d20}$$

[2]See Riddell (1981) for a seminal contribution here.

and

$$\frac{\varphi_h'(v)}{\varphi_w'(1-v)} = \frac{\Delta_h}{\Delta_w}, \tag{d21}$$

Note that the bargained outcome is *ex-post efficient* in the sense that the partners' marginal rates of substitution are equalized:

$$\frac{\varphi_w'(1-v)}{u_w'(c_w)} = \frac{\varphi_h'(v)}{u_h'(c_h)}. \tag{d22}$$

This relation summarizes the "ex-post contract curve" which is defined for a particular level of household income. Moreover, it is easy to see that the contract curve is monotonic: the higher is the husband's utility, the higher is c_h and v.

In any realized state of the world, there will thus be an ex-post bargained utility for each partner, which we denote by $\widetilde{U}_h(y_h, y_w)$ and $\widetilde{U}_w(y_h, y_w)$, along with actions $\widetilde{c}_i(y_h, y_w)$ and $\widetilde{v}(y_h, y_w)$. In a similar fashion each partner would associate each state of the world with a particular bargained indirect utility and actions.

For our comparative statics purposes we want to compare the outcome at two different states of the world that have the same total household income. Hence consider two states of the world $(\underline{y}, \overline{y})$ and $(\overline{y}, \underline{y})$ where $\overline{y} > \underline{y}$. Since total household income is the same at the two nodes, the utility possibility set is the same at the two nodes. However, comparative statics along the lines used above (or, noting that the shift from $(\underline{y}, \overline{y})$ to $(\overline{y}, \underline{y})$ is equivalent to an income redistribution) yields that

Lemma 6. *(Aizer, 2010) Consider two states of the world, $(\underline{y}, \overline{y})$ and $(\overline{y}, \underline{y})$ where $\overline{y} > \underline{y}$. Ex-post bargaining then implies that $\widetilde{U}_h(\underline{y}, \overline{y}) < \widetilde{U}_h(\overline{y}, \underline{y})$ and $\widetilde{U}_w(\underline{y}, \overline{y}) > \widetilde{U}_w(\overline{y}, \underline{y})$. Moreover, the ex-post negotiated violence level satisfies $\widetilde{v}(\underline{y}, \overline{y}) < \widetilde{v}(\overline{y}, \underline{y})$.*

We can now consider ex-ante bargaining with ex-post negotiations—i.e., bargaining once all uncertainty is resolved—as the fallback position.

D.4.4.2 The Ex-Ante Problem

Note that the resource allocation that the spouses would obtain through ex-post bargaining, $\{\widetilde{c}_h(y_h, y_w), \widetilde{c}_w(y_h, y_w), \widetilde{v}(y_h, y_w)\}$, is a feasible allocation according to Definition 1. Hence ex-post bargaining would generate an ex-ante expected utility for partner i

$$\widetilde{U}_i(F) = \sum_{y_h \in Y} \sum_{y_w \in Y} f_h(y_h) f_w(y_w) \widetilde{U}_i(y_h, y_w). \tag{d23}$$

Moreover, the expected utility profile $(\widetilde{U}_h(F), \widetilde{U}_w(F))$ is in the set T. However, noting that an allocation that would arise through ex-post bargaining is not ex-ante efficient, the expected utility profile $(\widetilde{U}_h(F), \widetilde{U}_w(F))$ is not a boundary element of T and hence it is Pareto dominated by some other element in T. Thus, both partners have an incentive to bargain for an ex-ante agreement, in this case with $\widetilde{U}_h(F)$ and $\widetilde{U}_w(F)$ as their respective fallback utilities.

In order to establish the result of interest, we need to verify that the husband's expected utility from ex-post bargaining is reduced from the shifting of probability defined in (d16) while that of the wife is increased. But this follows directly from Lemma 6. Hence by an analogous argument to the case with autarky as the threat point we obtain:

Proposition 3. *Suppose that the relevant threat point in the ex-ante bargaining process is ex-post bargaining. Then the shifting of probability in eq. (d16) leads to an decrease in the ex-ante bargained state-independent equilibrium level of violence* $v^* = v(y_h, y_w)$.

Bibliography

Aizer, A. (2010). The gender wage gap and domestic violence. *American Economic Review 100*, 1847–1859.

Aizer, A. and P. Dal Bó (2009). Love, hate and murder: Commitment devices in violent relationships. *Journal of Public Economics 93*(3), 412–428.

Albanesi, S. and A. Sahin (2013). The gender unemployment gap: Trend and cycle. Mimeo. Federal Reserve Bank of New York.

An, C.-B., R. Haveman, and B. Wolfe (1993). Teen out-of-wedlock births and welfare receipt: The role of childhood events and economic circumstances. *The Review of Economics and Statistics*, 195–208.

Anderberg, D. and Y. Zhu (2014). What a difference a term makes: the effect of educational attainment on marital outcomes in the UK. *Journal of Population Economics 27*(2), 387–419.

Angelucci, M. (2008). Love on the Rocks: Alcohol Abuse and Domestic Violence in Rural Mexico. *B.E. Journal of Economic Analysis and Policy: Contributions to Economic Analysis and Policy 8*, Article 43.

Angrist, J. and G. Imbens (1994). Identification and estimation of local average treatment effects. *Econometrica: Journal of the Econometric Society 62*, 467–475.

Angrist, J. and V. Lavy (2009). The effects of high stakes high school achievement awards: Evidence from a randomized trial. *The American Economic Review*, 1384–1414.

Angrist, J. D. and V. Lavy (1999). Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *The Quarterly Journal of Economics 114*(2), 533–575.

Armand, A. (2013). Who wears the trousers in the family? Intra-household resource control, subjective expectations and human capital investment. *University College London*.

Ashcraft, A., I. Fernández-Val, and K. Lang (2013). The Consequences of Teenage Childbearing: Consistent Estimates When Abortion Makes Miscarriage Non-random. *The Economic Journal 123*(571), 875–905.

Baird, S., E. Chirwa, C. McIntosh, and B. Özler (2010). The short-term impacts of a schooling conditional cash transfer program on the sexual behavior of young women. *Health Economics 19*(S1), 55–68.

Barrow, L. and C. E. Rouse (2013). Financial incentives and educational investment: The impact of performance-based scholarships on student time use. *NBER Working Paper no: 19351.*

Bartik, T. J. (1991). Who benefits from state and local economic development policies? W. E. Upjohn Institute for Employment Research. Kalamazoo, Michigan.

Becker, G. S. (1973). A Theory of Marriage: Part I. *The Journal of Political Economy*, 813–846.

Becker, G. S. (1974). A Theory of Marriage: Part II. *Journal of Political Economy 82*(2), pp. S11–S26.

Belot, M. and M. Francesconi (2013). Dating Preferences and Meeting Opportunities in Mate Choice Decisions. *Journal of Human Resources 48*(2), 474–508.

Bergstrom, T. and D. Lam (1989). The effects of cohort size on marriage markets in twentieth century Sweden. In T. Bengtsson (Ed.), *The Family, the Market, and the State in Industrialized countries. New perspectives on fertility in Britain. T. Bengtsson (Ed.)*, pp. 3. Oxford: Oxford University Press.

Bergstrom, T. C. and M. Bagnoli (1993). Courtship as a waiting game. *Journal of Political Economy*, 185–202.

Bergstrom, T. C., L. Blume, and H. Varian (1986). On the private provision of public goods. *Journal of Public Economics 29*, 25–49.

Berthelon, M. and D. Kruger (2010). Risky behavior among youth: Incapacitation effects of school on adolescent motherhood and crime in Chile. *Journal of Public Economics 95*, 41–53.

Bhaskar, V. (2012). Marriage market implications of the demographic transition. *Mimeo*.

Black, S., P. Devereux, and K. Salvanes (2008). Staying in the Classroom and out of the maternity ward? The effect of compulsory schooling laws on teenage births. *The Economic Journal 118*, 1025–1054.

Blanchard, O. J. and L. F. Katz. (1992). Regional evolutions. *Brookings Papers on Economic Activity 23*, 1–76.

Bloch, F. and V. Rao (2002). Terror as a bargaining instrument: A case study of dowry violence in rural India. *American Economic Review 92*, 1029–1043.

Braakmann, N. (2011). Female education and fertility–evidence from changes in British compulsory schooling laws. *Newcastle Discussion Papers in Economics 5*, 2011.

Brandt, L., A. Siow, and C. Vogel (2009). Large shocks and small changes in the marriage market for famine born cohorts in China. *IZA*, Discussion Paper Series, no. 4243.

Bronars, S. and J. Grogger (1994). The economic consequences of unwed motherhood: Using twin births as a natural experiment. *The American Economic Review 84*(5), 1141–1156.

Bronson, M. A. and M. Mazzocco (2012). Cohort Size and The Marriage Market: Explaining Nearly a Century of Changes in US Marriage Rates. *CCPR*, Working Paper.

Browning, M., P.-A. Chiappori, and V. Lechene (2010). Distributional effects in household models: Separate spheres and income pooling. *Economic Journal 120*, 786–799.

Campbell, D., A. Carruth, A. Dickerson, and F. Green (2007). Job insecurity and wages. *Economic Journal 117*, 544–566.

Card, D. and G. Dahl (2011). Family violence and football: The effect of unexpected emotional cues on violent behavior. *Quarterly Journal of Economics 126*(1), 103–143.

Chevalier, A. (2004). Parental education and child's education: A natural experiment. *IZA*, Discussion Paper Series, no. 1153.

Chevalier, A., C. Harmon, I. Walker, and Y. Zhu (2004). Does Education Raise Productivity, or Just Reflect it? *The Economics Journal 114*, F499–F517.

Chevalier, A. and T. Viitanen (2003). The long-run labour market consequences of teenage motherhood in Britain. *Journal of Population Economics 16*, 323–343.

Chioda, L., J. M. De Mello, and R. R. Soares (2012). Spillovers from conditional cash transfer programs: Bolsa Família and crime in urban Brazil. *IZA*, Discussion Paper Series, no. 6371.

Cho, I.-K. and D. M. Kreps (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics 102*(2), 179–221.

Chowdry, H., L. Dearden, and C. Emmerson (2008). Education maintenance allowance evaluation with administrative data: The impact of the EMA pilots on participation and attainment in post-compulsory education.

Clark, D., M. Geruso, and H. Royer (2014). The Impact of Education on Family Formation: Quasi-Experimental Evidence from the UK. *Mimeo*.

Clark, D. and H. Royer (2013). The Effect of Education on Adult Mortality and Health: Evidence from Britain. *American Economic Review 103*(6), 2087–2120.

Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy*, 957–976.

Crawford, C., L. Dearden, and C. Meghir (2010). When you are born matters: The impact of date of birth on educational outcomes in England. *IFS Working Papers*, No. 10, 06.

Danzer, N. and V. Lavy (2013). Parental Leave and Childrens Schooling Outcomes: Quasi-Experimental Evidence from a Large Parental Leave Reform. *NBER Working Paper no. 19452*.

Dearden, L., C. Emmerson, C. Frayne, and C. Meghir (2009). Conditional cash transfers and school dropout rates. *Journal of Human Resources 44*(4), 827–857.

Dearden, L. and A. Heath (1996). Income support and staying in school: What can we learn from Australia's AUSTUDY experiment? *Fiscal Studies 17*(4), 1–30.

Dee, T. S. (2001). Alcohol abuse and economic conditions: Evidence from repeated cross-sections of individual-level data. *Health Economics 10*, 257–270.

Devereux, P. J. and R. A. Hart (2010). Forced to be Rich? Returns to Compulsory Schooling in Britain. *The Economic Journal 120*(549), 1345–1364.

Díaz-Giménez, J. and E. Giolito (2013). Accounting for the timing of first marriage. *International Economic Review 54*(1), 135–158.

Dickson, M. and S. Smith (2011). What determines the return to education: An extra year or a hurdle cleared? *Economics of Education Review 30*(6), 1167–1176.

Doyle, O., C. Harmon, and I. Walker (2005). The impact of parental income and education on child health: Further evidence for england. *IZA*, Discussion Paper Series, no. 1832.

Ermisch, J. and D. Pevalin (2003). Who has a child as a teenager? *Institute for Social and Economic Research. Working Paper*.

Fan, J. and I. Gijbels (1996). Local Polynomial Modelling and Its Applications. *Monographs on Statistics and Applied Probability*, 66.

Farmer, A. and J. Tiefenthaler (1997). An economic analysis of domestic violence. *Review of Social Economy 55*, 337–358.

Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson (2006). Gender differences in mate selection: Evidence from a speed dating experiment. *The Quarterly Journal of Economics*, 673–697.

Fletcher, J. M. and B. L. Wolfe (2009). Education and Labor Market Consequences of Teenage Childbearing Evidence Using the Timing of Pregnancy Outcomes and Community Fixed Effects. *Journal of Human Resources 44*(2), 303–325.

Fort, M., N. Schneeweis, and R. Winter-Ebmer (2011). More Schooling, More Children: Compulsory Schooling Reforms and Fertility in Europe. *IZA Discussion Papers Working Paper No. 6015.*

Francesconi, M. (2008). Adult outcomes for children of teenage mothers. *Scandinavian Journal of Economics 110*(1), 93–117.

Friedberg, L. (1998). Did Unilateral Divorce Raise Divorce Rates? Evidence from Panel Data. *The American Economic Review 88*(3), 608–627.

Fryer, R. G. (2011). Financial Incentives and Student Achievement: Evidence from Randomized Trials. *The Quarterly Journal of Economics 126*(4), 1755–1798.

Galindo-Rueda, F. (2003). The intergenerational effect of parental schooling: Evidence from the British 1947 school leaving age reform. *Centre for Economic Performance, London School of Economics, mimeo*.

Gelman, A. and G. Imbens (2014). Why high-order polynomials should not be used in regression discontinuity designs. *NBER*, Working Paper no. 20405.

Grabill, W. R. and L. J. Cho (1965). Methodology for the measurement of current fertility from population data on young children. *Demography 2*(1), 50–73.

Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2014). Marry your like: Assortative mating and income inequality. *The American Economic Review 104*(5), 348–353.

Grönqvist, H. and C. Hall (2013). Education policy and early fertility: Lessons from an expansion of upper secondary schooling. *Economics of Education Review 37*, 13–33.

Gutiérrez, J. P., S. Bautista, P. Gertler, M. Hernández, and S. M. Bertozzi (2004). External evaluation of the impact of the human development program 'Oportunidades'. *Mexico City*.

Hahn, J., P. Todd, and W. van der Klaauw (1999). Evaluating the effect of an antidiscrimination law using a regression-discontinuity design. *NBER Working Paper no. 7131.*

Hahn, J., P. Todd, and W. van der Klaauw (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica: Journal of the Econometric Society 69*(1), 201–209.

Handa, S., C. T. Halpern, A. Pettifor, and H. Thirumurthy (2014). The Government of Kenya's Cash Transfer Program Reduces the Risk of Sexual Debut among Young People Age 15-25. *PloS one 9*(1).

Harmon, C. and I. Walker (1995). Estimates of the economic return to schooling for the United Kingdom. *The American Economic Review 85*(5), 1278–1286.

Harris, K., G. Duncan, and J. Boisjoly (2002). Evaluating the role of nothing to lose attitudes on risky behavior in adolescence. *Social Forces 3*, 1005–1039.

Holmlund, H. (2008). Intergenerational mobility and assortative mating: Effects of an educational reform. *Centre for the Economics of Education, LSE*.

Hotz, V., S. McElroy, and S. Sanders (2005). Teenage childbearing and its life cycle consequences: Exploiting a natural experiment. *Journal of Human Resources 40*, 683–715.

Imbens, G. W. and T. Lemieux (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics 142*(2), 615–635.

Iyengar, R. (2009). Does the certainty of arrest reduce domestic violence? Evidence from mandatory and recommended arrest laws. *Journal of Public Economics 93*(1), 85–98.

Jacob, B. A. and L. Lefgren (2003). Are Idle Hands the Devil's Workshop? Incapacitation, Concentration, and Juvenile Crime. *American Economic Review 5*, 1560–1577.

Kearney, M. S. and P. B. Levine (2012). Why is the teen birth rate in the United States so high and why does it matter? *NBER Working Paper no. 17965.*

Kiernan, K. E. (1997). Becoming a young parent: A longitudinal study of associated factors. *British Journal of Sociology*, 406–428.

Klepinger, D., S. Lundberg, and R. Plotnick (1995). Adolescent fertility and the educational attainment of young women. *Family Planning Perspectives 27*, 23–28.

Lee, D. S. (2001). The Electoral Advantage to Incumbency and Voters' Valuation of Politicians' Experience: A Regression Discontinuity Analysis of Elections to the US. *NBER Working Paper no. 8441*.

Lee, D. S. and D. Card (2008). Regression discontinuity inference with specification error. *Journal of Econometrics 142*(2), 655–674.

Lee, D. S. and T. Lemieux (2010). Regression Discontinuity Designs in Economics. *The Journal of Economic Literature 48*(2), 281–355.

León, A. (2004). The Effect of Education on Fertility: Evidence from Compulsory Schooling Laws. *Mimeo*.

Leuven, E., H. Oosterbeek, and B. van der Klaauw (2010). The effect of financial rewards on students'achievement: Evidence from a randomized experiment. *Journal of the European Economic Association 8*(6), 1243–1265.

Lindeboom, M., A. Llena-Nozal, and B. van der Klaauw (2009). Parental education and child health: Evidence from a schooling reform. *Journal of Health Economics 28*(1), 109–131.

Lochner, L. and E. Moretti (2004). The Effect of Education on Crime: Evidence from Prison Inmates, Arrests, and Self-Reports. *The American Economic Review 94*(1), pp. 155–189.

Ludwig, J. and D. L. Miller (2007). Does Head Start improve children's life chances? Evidence from a regression discontinuity design. *The Quarterly Journal of Economics 122*(1), 159–208.

Machin, S., O. Marie, and S. Vujić (2011). The crime reducing effect of education. *The Economic Journal 121*(552), 463–484.

Mansour, H. and T. McKinnish (2013). Who Marries Differently-Aged Spouses? Ability, Education, Occupation, Earnings, and Appearance. *Review of Economics and Statistics 96*(0), 577–580.

McCrary, J. and H. Royer (2011). The effect of female education on fertility and infant health: Evidence from school entry policies using exact date of birth. *The American Economic Review 101*(1), 158–195.

Middleton, S., K. Perren, S. Maguire, J. Rennison, E. Battistin, C. Emmerson, and E. Fitzsimmons (2005). Evaluation of education allowance pilots: young people aged 16 to 19 years. final report of the quantitative evaluation. *Copyright: Queens Printer and Controller of HMSO*.

Monstad, K., C. Propper, and K. Salvanes (2008). Education and Fertility: Evidence from a Natural Experiment. *Scandinavian Journal of Economics 110*, 827–852.

Moore, K. and L. Waite (1978). The impact of an early first birth on young women's educational attainment. *Social Forces 56*, 845–865.

Murphy, M. and A. Berrington (1993). Constructing period parity progression ratios from household survey data. In *New perspectives on fertility in Britain. Studies on medical and population subjects, M. Ní Bhrolchaín (Ed.)*, pp. 17–32. University of Chicago Press.

Office of National Statistics (2002). *Social Trends*. Palgrave MacMillan.

Office of National Statistics (2008). *Social trends*. Palgrave MacMillan.

Office of National Statistics (2010a). *Child Mortality Statistics*. UK Statistics Authority.

Office of National Statistics (2010b). *Cohort Fertility, England & Wales*. UK Statistics Authority.

Oreopoulos, P. (2006). Estimating average and local average treatment effects of education when compulsory schooling laws really matter. *The American Economic Review*, 152–175.

Oreopoulos, P. (2007). Do dropouts drop out too soon? Wealth, health and happiness from compulsory schooling. *Journal of public Economics 91*(11), 2213–2229.

Öster, A. and J. Agell (2010). Crime and unemployment in turbulent times. *Journal of the European Economic Association 5*, 752–775.

Paniagua, M. N. and I. Walker (2012). The Impact of Teenage Motherhood on the Education and Fertility of their Children: Evidence for Europe. *IZA Discussion Papers Working Paper No. 6995*.

Pollak, R. A. (2004). An intergenerational model of domestic violence. *Journal of Population Economics 17*, 311–329.

Raphael, S. and R. Winter-Ebmer (2001). Identifying the effect of unemployment on crime. *Journal of Law and Economics 44*, 259–283.

Rawlings, L. B. and G. M. Rubio (2005). Evaluating the impact of conditional cash transfer programs. *The World Bank Research Observer 20*(1), 29–55.

Rees, G., H. Williamson, and D. Istance (1996). Status Zero: A study of jobless school-leavers in South Wales. *Research Papers in Education 11*(2), 219–235.

Riccio, J., N. Dechausay, D. Greenberg, C. Miller, Z. Rucks, and N. Verma (2010). Toward reduced poverty across generations: Early findings from New York City's conditional cash transfer program. *MDRC, March*.

Riddell, W. C. (1981). Bargaining under uncertainty. *American Economic Review 71*(4), 579–590.

Sabates, R. and L. Feinstein (2007). Effects of government initiatives on youth crime. *Oxford Economic Papers*, 462–483.

Schmidt, S. (1999). Long-run trends in workers beliefs about their own job security: Evidence from the General Social Survey. *Journal of Labor Economics 17*, 127–141.

Schwartz, C. R. and R. D. Mare (2005). Trends in educational assortative marriage from 1940 to 2003. *Demography 42*(4), 621–646.

Silles, M. (2011). The effect of schooling on teenage childbearing: Evidence using changes in compulsory education laws. *Journal of Population Economics 24*, 761–777.

Skirbekk, V., H.-P. Kohler, and A. Prskawetz (2004). Birth month, school graduation, and the timing of births and marriages. *Demography 41*(3), 547–568.

Smith, K., S. Osborne, I. Lau, and A. Britton (2012). Homicides, firearm offences and intimate violence 2010/11. *Home Office Statistical Bulletin*.

Social Exclusion Unit (2001). Preventing social exclusion. *Crown Copyright. HM Cabinet Office*.

Stillwell, J., P. Boden, and P. Rees (1990). Trends in internal net migration in the UK: 1975 to 1986. *Area 22.1*, 57–65.

Tauchen, H. V., A. D. Witte, and S. K. Long (1991). Violence in the family: A non-random affair. *International Economic Review 32*, 491–511.

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

Trochim, W. M. (1984). Research Design for Program Evaluation: The Regression-Discontinuity Approach. *Sage Newbury Park, CA.*.

Warr, P. G. (1983). The private provision of a pure public good is independent of the distribution of income. *Economics Letters 13*, 207–211.

Wolfers, J. (2003). Did unilateral divorce laws raise divorce rates? A reconciliation and new results. *NBER Working Paper, No. 10014*.

Wood, J. and M. Weinstein (1988). A Model of Age-Specific Fecundability. *Population Studies 42*(1), 85–113.

Woodin, T., G. McCulloch, and S. Cowan (2013). Secondary Education and the Raising of the School-Leaving Age: Coming of Age? *Palgrave Macmillan*.