The Impact of Labour Market Conditions on Education Decisions

by Eric K. Mackay*

An essay submitted to the Department of Economics
in partial fulfillment of the requirements for
the degree of Master of Arts

Queen's University
Kingston, Ontario, Canada
August 2014

*The author would like to thank Vincent Pohl for his supervision of this MA Essay and for his suggestions, which helped shape the direction it has taken, Christopher Ferrall for his assistance during the early stages of researching this paper, and Shari Eli and Kory Kroft for their helpful suggestions when the idea for this paper was still in its infancy.

Abstract

Economic models of education decisions predict that increases in the expected return to education and reductions in the associated opportunity cost should encourage greater enrolment. Labour market conditions play a major role in driving these costs and benefits. This paper uses American Community Survey data from the years 2005 to 2011 in order to examine the impact of two major labour market factors—unemployment rates and earnings—on the enrolment decisions of 19 to 21-year-olds and 23 to 35-year-olds. Due to the plausible exogeneity of unemployment rates at the individual level, unemployment effects are estimated using probit regressions. Since enrolment decisions are predicted to depend on relative instead of absolute earnings by education, plausibly exogenous variation in the value of local oil reserves in the U.S. South is first shown to influence earnings differentials by education level. A measure of local exposure to oil shocks is subsequently used to instrument the differential between average high school and college graduate earnings in order to arrive at causal estimates of the relationship between the earnings premium and college enrolment.

Contents

Introduction

Economists have had a long-standing interest in determining the factors that influence individuals' school enrolment decisions. It has long been recognized that investments in education play an important role in both the labour market outcomes of individuals and in economic growth more broadly. Since the formulation of a basic economic model of education investment decisions by economists such as Gary Becker, much research has been undertaken to decompose trends in enrolment, test the performance of the basic model and its extensions, and estimate the impacts of policies and shocks on individual education decisions.

The goal of this paper is to provide further insight into the impacts of labour market conditions on the post-secondary enrolment decisions of college-aged individuals and the re-enrolment decisions of adults aged 23 to 35. This paper will examine the impacts of changes in two labour market factors in particular: the local age-specific unemployment rate, and average earnings by education level. This paper aims to make several contributions to the existing literature. Firstly, the unemployment-enrolment analysis section will attempt to differentiate between unemployment effects and recessionary effects more broadly by making use of recent data that includes the 2008 "Great Recession" period. Secondly, this paper will extend the analysis of the enrolment effects of earnings premia by education—which has tended to focus on high school enrolment—to the post-secondary enrolment decision. Lastly, this paper will attempt to provide a fairly broad review of earlier findings on the unemployment-enrolment and earnings premium-enrolment relationships.

This paper aims to estimate the causal effects of labour market conditions on enrolment decisions. It has been argued by van den Berg et al. (2006) that macroeconomic conditions can serve as a sort of "instrument" in micro-level analyses due to their exogeneity at the individual level. Following this reasoning, individual enrolment probabilities will be estimated as a function of local unemployment rates.

In order to estimate the enrolment effects of changes in wage premia by education level, this paper will exploit plausibly exogenous variation in the oil price and local oil reserves using an instrumental variables approach that closely follows Black et al. (2005) and Acemoglu et al. (2013).

Section 1 of this paper will provide some background and a broad review of the literature on the enrolment impacts of labour market conditions as well as a brief overview of some papers using quasi-experimental methods that are relevant to the analysis to follow. Section 2 will outline the empirical methodology used in this paper. Section 3 will provide a description of the data used in the analysis. Section 4 will explain the results, and will be followed by a conclusion.

1 Background and Literature Review

1.1 Background

Many of the foundations of modern economic models of human capital investment have their origins in the work of Gary Becker during the 1960s. Becker (1962) argued that the decision to invest in education must hinge on a comparison of the net present value of the expected return to further education against the net present value of the associated costs. An increase in the wage differential or decrease in the risk regarding future returns would raise the expected return to education, while a decrease in tuition fees or interest rates on borrowing would lower the costs. An individual's decision to enrol in school also imposes an implicit cost by forcing them to forgo opportunities to work full-time. Additionally, Becker notes that the utility or disutility an individual receives from learning may also be a factor in the enrolment decision.

Light (1996) argues against the prevailing view in early economic models of education investment that decisions are made, by and large, prior to entry into the labour market. She finds evidence that a substantial fraction of individuals in the

U.S. who leave school after graduating from high school later return to school. She argues that the decision to enrol should be viewed as a continuous weighing of the costs and benefits, which means that individuals might drop out and later return due to changes in labour market conditions or borrowing constraints.

Card and Lemieux (2000) provide a broad overview of the basic economic model of human capital investment with a number of extensions in their attempt to identify the drivers of U.S. enrolment trends from 1968 to 1996. Some of the extensions they outline include the importance of family background and entry cohort size (in the case of college enrolment) as potential explanatory factors. They also emphasize that some attempts to explain aggregate enrolment trends may suffer due to failures to account for local variation in education costs and labour market conditions. They also explain that booms, liquidity/credit constraints, or excessive discounting could influence opportunity costs.

1.2 The Impact of Unemployment Rates on Enrolment

There have been a large number of papers examining the relationship between unemployment rates and school enrolment. Edwards (1976) examines the relationship between high school enrolment and the unemployment rate among 16 to 17-year-olds in the U.S. between 1947 and 1974. She finds a significant positive correlation for females and a negative one for non-whites. Gustman and Steinmeier (1981) examine the impact of both wages and unemployment on the joint labour force participation-enrolment decision for youths using cross-sectional data from the 1976 U.S. Survey of Income and Education with a discrete multivariate analysis. They find statistically significant lower enrolment rates for males in metropolitan areas with low unemployment and the reverse for non-white females.

Betts and McFarland (1995) investigate the determinants of community college enrolment using a variation of the standard economic model of education decisions, which is more suited to the incentives facing mature students than typical four-year college models. They find that community college enrolment in the U.S. during the 1969 to 1985 period increased substantially during periods of high unemployment. Dellas and Sakellaris (2003) develop a partial equilibrium model incorporating the core elements of the Becker model and allowing for credit and financing elements, and for heterogeneity in the cost of acquiring education and in initial asset endowments. The model arrives at similar conclusions with regards to the effects of changes in wage premia by education level and education costs. They find that decreases in growth and increases in unemployment both correspond with increased college enrolment using U.S. Current Population Survey data from 1968-1988. Dellas and Koubi (2003) formulate a similar partial equilibrium model and analyse several U.S. time-series on school enrolment rates from various sources in order to determine the unemployment-enrolment relationship for several age groups. They find a statistically significant relationship between the overall unemployment and enrolment rates, as well as the unemployment and enrolment rates for several age groups. Card and Lemieux (2000), in the process of determining the drivers of U.S. enrolment trends over the 1968 to 1996 period, find modest positive relationships between unemployment rates and enrolment among groups of 15 to 16-year-olds and 17-year-olds. They also find negative point estimates for the 18-year-old and 19 to 21-year-old categories, which they note could be explained by the effects of unemployment on parental income or by borrowing constraints among college-age cohorts, although none of these estimates are statistically significant. Light (1996) estimates a proportional hazard model which examines the duration of non-enrolment among men who completed high school in the 1979 to 1991 years of the National Longitudinal Survey of Youths. She finds that higher unemployment rates correspond with a greater likelihood that an individual re-enrols.

Several papers have been produced in the past few years that attempt to address the effects of unemployment on school enrolment. Clark (2011) uses U.K. panel data on 8 regions over the period of 1975-2005 which include specific unemployment measures for youths, in contrast to most of the earlier studies done using U.S. data. He finds that youth unemployment and school enrolment are positively related. Using a directed search model, Guo (2013) investigates the impact of recessions on schooling decisions, wages, and employment. She fits her model using data on U.S. cohorts that reached the age of 16 between 1933 and 2012 with a simulated method of moments approach. She uses her model to simulate a number of shocks which yield a countercyclical pattern in school enrolment. Furthermore, she performs a counterfactual analysis using her model which suggests that the 1981-1983 recession increased average years of schooling.

1.3 The Impact of Unskilled Labour Wages and Minimum Wages on Enrolment

A handful of papers have been written on the effects of unskilled labour wages on the enrolment decisions of youths. In their investigation of the joint enrolment-labour force participation decision for youths using 1976 U.S. data, Gustman and Steinmeier (1981) find lower enrolment and higher labour force participation in metropolitan areas with higher relative youth wages. However, these results are found to be sensitive to adjustments to account for full- and part-time work status. Black et al. (2005) find evidence that the Appalachian coal boom in the 1970s led to a relative increase in low-skilled wages in coal-rich counties and that the subsequent bust in the 1980s reversed this trend. Using this variation, they find that increases in low-skilled wages resulted in substantial decreases in high school enrolment. Emery et al. (2011) examine the impact of the OPEC oil boom and subsequent bust on high school enrolment and long-run educational attainment in Alberta. They apply a difference-in-differences approach using the International Adult Literacy Survey (IALS) that compares the education levels of pre-boom (born 1953-1956) and during-boom (born 1961-1964) cohorts in Alberta against the rest of Canada. They find that the male boom cohort had higher rates of non-university post-secondary education completion. They also apply a synthetic cohort methodology using data from the Canadian Census. Their results provide some evidence that boom cohorts might be more likely to delay post-secondary education instead of forgoing it altogether, whereas those born during the bust fared worse in terms of post-secondary attainment. Light (1996) also finds some evidence of the impact of wages on school re-enrolment. Her results suggest that individuals in the workforce with higher wages have a lower chance of re-enrolling in school at a later date.

Some of the existing research on the role of wages in the education decisions of youths has focused on the effects of minimum wages. Neumark and Wascher (1995) use May CPS data from 1979-1992 with individuals matched across 2 periods in order to estimate the impact of minimum wages relative to average wages by state on the probability that a 16 to 19-year-old is in school and working, in school and not working, or neither in school nor working. They find that higher minimum wages correspond with a decreased probability of being in school and, when conditioning on individuals' work/school statuses, they find some evidence that higher minimum wages correspond with a lower probability of remaining in school and a higher probability of leaving school. Chaplin et al. (2003) find a negative relationship in the U.S. between high school continuation rates and state minimum wages using the Common Core of Data from 1989 to 1997. They find evidence that the effect is largest during the grade 9 to grade 10 transition.

1.4 Quasi-Experiments Exploiting Plausibly Exogenous Variation in Earnings

Several recent papers have used exogenous variation in oil and coal prices during the 1970's and 1980's to instrument earnings in various U.S. local geographic areas. Black et al. (2005) use variation in the value of coal reserves in the Appalachian counties to determine the causal influence of wages on high school drop-out decisions. They find that the wage differential between high school graduates and drop-outs varies by local coal intensity. So as to estimate the causal effect of the wage differential on drop-out rates, they instrument earnings using the interaction between the amount of local coal reserves and the coal price. Black et al. (2002) use a similar approach involving variation in the coal price and reserves from 1970-1993 across the Appalachian counties in order to estimate the relationship between local earnings and the use of Disability Insurance (DI) and Supplemental Security Income (SSI). Additionally, they instrument county earnings using the fraction of employment in primary metals in 1970 to exploit the plausibly exogenous decline in the industry to estimate these effects. Acemoglu et al. (2013) exploit the variation in local oil intensity and oil prices in the U.S. South over the 1970 to 1990 period to estimate the causal effect of incomes on health expenditures. Since income is expected to influence health spending and healthier workers are expected to generate more income, they attempt to disentangle the effect of incomes on health expenditures by using the fact that oil booms and busts should influence incomes independently of productivity-improving investments in health. They instrument incomes at the economic subregion level using the interaction between the oil price and local total oil reserves since discovery in order to estimate the effects of income on health expenditures.

There are also papers focusing on other countries that take advantage of the influence of oil booms and busts on labour market conditions. Emery et al. (2011) compare high school enrolment and long-run educational attainment in oil-rich Alberta against the rest of Canada among high school-aged cohorts before and after the 1970's oil boom in a difference-in-differences approach. Loken (2007) exploits exogenous variation in family incomes resulting from the Norwegian oil boom in the 1970's and 1980's. She restricts her sample to two counties, one with direct exposure to the oil boom and the other with no direct exposure, and instruments average family income over the 1973 to 1988 period using a dummy variable specifying whether an individual resided in the county with significant oil exposure. (It should be noted, however, that instrumenting family income using a dummy in this manner is unlikely to yield superior results to a difference-in-differences alternative.) Marchand (2012) divides the Western Canadian Census divisions into a treatment group with significant exposure to oil prices (based on energy revenues) and a comparison group. He then uses a difference-in-differences approach that compares employment and earnings outcomes of the two groups during the 1971-1981 and 1996-2006 energy booms and the 1981-1991 energy bust.

2 Empirical Methodology

The empirical analysis to follow will attempt to test two hypotheses: that increases in unemployment result in increased school enrolment, and that relative increases in the wages of low-skilled workers reduce post-secondary enrolment among college-aged youths. Both of these hypotheses readily follow from the standard economic model of human capital investment since a rise in the unemployment rate will decrease the opportunity cost of school (as the prospects of the alternative of working deteriorate) and a decrease in the wage gap between workers with different levels of human capital will lower the net present value of the expected returns to further education.

It is arguable that the coefficient estimates from simple probit regressions of individual enrolment status on the local unemployment rate can be interpreted as the causal effect of increases in the probability of enrolment in response to changes in the unemployment rate. This is because the unemployment rate is exogenous at the individual level (see van den Berg et al. (2006))—that is, individuals take the unemployment rate as given and their enrolment decisions have no meaningful impact on the overall unemployment rate. However, regions with consistently lower levels of human capital investment may tend to suffer from chronically higher rates of unemployment which could introduce confounds into a naive probit regression. Furthermore, individuals in high unemployment regions who are not in school might be more likely to move to other regions in search of jobs which could exaggerate any positive effects that the unemployment rate may have on enrolment. Nonetheless, these concerns can be addressed through the inclusion a proxy for migration and of regional dummy variables to control for pre-existing trends. Equation 1 outlines the probit regression that will be run.

$$Pr(enrolled_i = 1) = \Phi(\beta_0 + \beta_1 unrate_{jt} + X_i\beta + \epsilon_i) \tag{1}$$

where, for an individual i residing in public use microdata area (PUMA) j and surveyed in year t, $enrolled_i$ is a dummy variable which equals 1 if the individual is enrolled in school, $unrate_{jt}$ is the local age-specific unemployment rate for PUMA j in year t, X_i is a vector of controls including PUMA of residence dummies set equal to 1 if the individual resides in the respective PUMA and a proxy for the local migration rate in the individual's PUMA of residence, and ϵ_i is an error term. This probit regression will be run for both the probability of post-secondary enrolment among 19 to 21-year-olds and for school enrolment of all types among 23 to 35-year-olds.

In contrast to the proposed analysis of the effects of unemployment on enrolment, measuring the impacts of changes in earnings on enrolment presents a number of problems. Firstly, in the earlier probit regressions, all individuals in a given region faced the same unemployment rate whereas wages are only observed for those who are employed. This presents difficulties since analogous probit regressions for the enrolment-earnings relationship would require knowledge of the hypothetical wage that an individual who is currently enrolled would face were they to enter the labour market looking for full-time work. Secondly, while local average wages could be used as a plausibly exogenous proxy for the hypothetical wage that an individual would face in the market, a procedure analogous to the one suggested for estimating unemployment effects would still be problematic. The reason for this is that economic models of human capital investment predict that the enrolment decision should depend on the wage differential by education level, not the absolute wage that an individual might face. Lastly, there is no reason to expect that the earnings differential by education level is likely to be sufficiently variable in order to facilitate the estimation of its effects on enrolment.

In order to estimate the causal effect of changes in the earnings differential–the return to education–on the enrolment decisions of college-aged individuals, this paper will exploit plausibly exogenous variation in the oil price and local oil intensity using an instrumental variables approach. First it will be shown that the wage differential varies with local exposure to oil shocks. This follows from the fact that natural resource booms are anticipated to disproportionately benefit less skilled workers and is supported by the findings of Black et al. (2005) in their study of the effects of the coal boom and bust over the 1970's and 1980's on high school enrolment. In order to test this, the following regression will be estimated for individuals aged 25 to 30:

$$income_i = \beta_0 + \beta_1(z_{jt} \times hsoe_i) + \beta_2(z_{jt} \times spse_i) + \beta_3(z_{jt} \times psom_i) + \beta_4 hsoe_i + \beta_5 spse_i + \beta_6 psom_i + \beta_7 z_{jt} + X_i\beta + \epsilon_i \qquad (2)$$

where, for an individual i residing in PUMA j and surveyed in year t, $hsoe_i$, $spse_i$, and $psom_i$ are dummies set equal to 1 if the individual's highest level of education is high school or an equivalent, some post-secondary education, or post-secondary or more, respectively, z_{jt} is an instrument constructed by interacting the oil reserves in PUMA j with the oil price in year t, X_i is a vector of controls, including potential experience, and ϵ_i is an error term. In this regression, the omitted educational attainment category is high school drop-outs.

After determining this relationship, the average post-secondary earnings premium among 25 to 30-year-olds will be instrumented using a measure of local oil intensity interacted with the oil price in an instrumental variables regression of the local enrolment rate on the local post-secondary earnings premium. The procedure used is similar in many respects to those of Black et al. (2005) and Acemoglu et al. (2013). Equations 3 and 4 outline the proposed second and first stages of the IV regression, respectively:

$$enrrate_{jt} = \beta_0 + \beta_1 \widehat{premium}_{jt} + X_{jt}\beta + \epsilon_{jt} \qquad (3)$$

where

$$premium_{jt} = \gamma_0 + \gamma_1 z_{jt} + X_{jt}\gamma + \upsilon_{jt} \tag{4}$$

and j and t are indexes for the public use microdata area (PUMA) and year, respectively, $enrrate_{jt}$ is the school enrolment rate, $premium_{jt}$ is a measure of the annual average wage/salary income differential for a specified amount of additional education, $\widehat{premium}_{jt}$ is the fitted value of $premium_{jt}$ obtained from running the IV first stage regression outlined in equation (4), z_{jt} is the instrument, X_{jt} is a set of control variables including a proxy for the migration rate, and ϵ_{jt} and υ_{jt} are error terms.

The measure of local oil intensity that will be used to construct the instrument is the value of oil reserves at the PUMA level. The instrument itself is the product of the PUMA-level oil reserves and the oil price and can be interpreted as the value of local oil reserves. The instrumental variables methodology crucially depends on the exogeneity of the instrument. One potential threat to this exogeneity requirement is the fact that oil booms might attract workers who are less likely to decide to enrol in school to move to oil-rich PUMAs in search of higher wages. This could lead to biased estimates of the enrolment-earnings premium relationship. As such, a proxy for migration—the average number of residents who moved in the past year—will be included among the controls. Since both the oil price and initial endowments of oil across PUMA's are almost certainly free from the influence of highly localized events at the PUMA-level, and since, after controlling for migration, the value of local oil reserves is only expected to influence enrolment through its impacts on local labour market conditions, it should function as a valid instrument.

It should be noted that, in contrast to Acemoglu et al. (2013), it was only possible to obtain recent estimates of oil reserves instead of estimates of the initial endowment. This could introduce some bias into the IV regression if there is a significant correlation between past oil extraction intensity (which would effect current reserves) and key omitted variables. Nonetheless, this would likely result in underestimates of any effects of earnings premia on enrolment since PUMAs that were the most aggressive in exploiting their oil reserves would be at a greater risk of being treated as having less exposure to oil in the event that they used up much of their oil reserves. It is also possible that even the estimated initial endowment of oil reserves might not be perfectly exogenous since the patterns of energy exploration on which the oil reserve estimates depend may be correlated with omitted variables. However, any such risks are likely to be very small.

3 Data

The dataset used in this study is a 1% sample of the American Community Survey (ACS) for the years 2005 through 2011 obtained from the University of Minnesota's Integrated Public Use Microdata Series (IPUMS-USA) database. The public use version of the ACS collects extensive information on school enrolment, educational attainment, earnings, and labour force status, among other factors. The survey also provides a substantial amount of geographic detail for respondents which was crucial for implementing the instrumental variables approach pursued in this paper. The smallest geographic unit is the Public Use Microdata Area (PUMA) which combines counties with small populations together until a sufficiently large population is obtained. The IV regression performed in this paper also required the creation of a panel of PUMA-level aggregates from the ACS microdata. The aggregates were generated using the ACS person weights and include various PUMA-level school enrolment rates, average wages, and unemployment rates.

The working dataset was narrowed down to the years 2005 through 2011 so that the PUMA definitions would be consistent across years. The U.S. top 100 oil fields by reserves in 2009 are found in the Southern States, the Mid-West, California, and Alaska. It was decided to focus largely on the Southern states to avoid problems introduced by the paucity of PUMAs in most Mid-Western states due to their low population densities, and to improve cross-state comparability. States lacking significant onshore oil reserves were removed from the sample. This allows for the exploitation of both within-state and cross-state variation in oil reserves and reduces the risk of any IV estimates being driven solely by cross-state variation that is unrelated to oil reserves. The final sample consists of six states—Texas, New Mexico, Oklahoma, Louisiana, Mississippi, and Colorado—each of which contain at least one of the 2009 top 100 oil fields[1].

In order to construct the value of oil reserves instrument it was necessary to obtain an estimate of the oil reserves in each of the PUMAs in the sample. Reserve data was obtained from the U.S. Energy Information Administration's Top 100 Oil and Gas Fields of 2009 supplement to their 2009 U.S. Crude Oil, Natural Gas, and Natural Gas Liquids Proved Reserves report [2]. Classification of oil fields into their respective PUMAs was done with reference to the IPUMS-USA PUMA documentation [3] in combination with a variety of sources, since the EIA oil field information does not specify the county/counties in which the oil fields are located. These sources included, but were not necessarily limited to, the websites of geological associations, oil and gas companies, state historical societies, and mapping/geolocation services. Some judgement was necessary in order to determine whether or not the oil field extended into

[1] It is worth noting that school attendance is compulsory until age 18 in Texas, New Mexico, Oklahoma, and Louisiana, and age 17 in Mississippi, and Colorado according to the National Center for Education Statistics (`http://nces.ed.gov/programs/statereform/tab5_1.asp`). These high legal drop-out ages were the main factor behind the decision not to include an analysis of high school enrolment decisions in this paper.

[2] `http://www.eia.gov/naturalgas/crudeoilreserves/archive/2009/pdf/top100fields.pdf`

[3] `https://usa.ipums.org/usa/volii/2000pumas.shtml`

neighbouring counties although these decisions rarely affected the PUMA-level classifications since there was little ambiguity as to the primary county/counties containing the oil field and since neighbouring counties were frequently part of the same PUMAs due to the level of county aggregation in the sparsely populated oil-rich counties. In total, there are 21 oil-rich PUMAs in the sample relative to a total of 281 PUMAs which appear across all seven years of the sample.

The value of oil reserves instrument was created by interacting the 2009 PUMA level oil reserves with the real annual average of WTI oil spot prices at Cushing, OK from Jan. 1, 2005 to Dec. 31, 2011. This data was obtained from the Federal Reserve Economic Data (FRED) at the Federal Reserve Bank of St. Louis and originated from Dow Jones & Company via the Wall Street Journal.

PUMA-level statistics such as the unemployment rate, enrolment rates, and average incomes calculated using a 1% sample of the ACS are likely to suffer from a modest amount of sampling error. The median number of respondents observed per year in each PUMA was 1,341 with a minimum of 539 and a maximum of 3,696. For some subgroups in some PUMAs this results in summary statistics that are computed using reasonably small numbers of respondents. For example, in 2007, PUMA 700 in Colorado had 169 respondents between the ages of 23 and 35, all of whom were either employed or not in the labour force. This resulted in an estimated unemployment rate of 0% which is, without a doubt, an underestimate. However, such an estimate is unsurprising given that the unemployment rates in this PUMA during the preceding and following years were found to be 0.9% and 0.3%, respectively. As a result of such imprecisions, estimated standard errors in regressions using PUMA-level statistics are likely to underestimate the true standard errors somewhat. This should be borne in mind when interpreting the results of this paper.

The analysis to follow makes use of PUMA-level unemployment rates computed using employment status information collected as part of the ACS. The Bureau of

Labor Statistics (BLS) notes that the ACS unemployment estimates tend to be higher than the local area unemployment statistics (LAUS) obtained from the Current Population Survey (CPS), likely as a result of the less detailed questions regarding job search activity in the ACS and due to differences in the reference period[4]. The BLS also notes that the labour force questions in the ACS were refined somewhat in 2008. While it is recognized that the ACS unemployment estimates are likely to be of lower quality than the CPS-LAUS, it was decided to use the ACS measure in order to ensure geographic consistency at the PUMA level. At the very least, the ACS unemployment rate can be viewed as a reasonable proxy for the CPS-LAUS measure.

The income and oil price measures used in this paper are all adjusted to 1999 dollars using the national CPI-U index included in the ACS dataset (which, IPUMS notes, originates from the Bureau of Labor Statistics). The Census Bureau also provides an adjustment factor to account for the fact that individuals are surveyed at different times over the course of the year and therefore have different reference frames for reporting their income. It was decided not to apply this adjustment based on the recommendation of IPUMS-USA due to flaws in its construction.

Table 1 reports summary statistics for the aggregate level ACS dataset for both oil-rich and oil-poor PUMAs. The table reveals that enrolment rates for both 19 to 21-year-olds and 23 to 35-year-olds are noticeably lower in oil-rich PUMAs (by about 6 and 2.5 percentage points, respectively). Unemployment rates are also lower in the oil-rich PUMAs with a difference of 1.35 percentage points for 19 to 21-year-olds and 0.74 percentage points for 23 to 35-year-olds. Average incomes for high school dropouts and high school graduates are a little over $200 dollars higher in oil-rich PUMAs compared to oil-poor ones, while average incomes for post-secondary graduates are over $2,000 lower. Moving rates, while quite high for both age groups in comparison to the overall population average of about 15% (not reported), are about 3 percentage

[4] `http://www.bls.gov/lau/acsqa.htm`

points higher for 19 to 21-year-olds in oil-rich PUMAs.

Table 1: Summary Statistics for 2005-2011 PUMA Education, Unemployment, and Earnings

	All PUMAs		Oil-Rich PUMAs		Oil-Poor PUMAs	
Variables	Obs	Mean (Std.Dev.)	Obs	Mean (Std.Dev.)	Obs	Mean (Std.Dev.)
19 to 21-year-old undergraduate enrol.	1969	43.19% (16.21%)	148	37.47% (14.25%)	1821	43.65% (16.27%)
23 to 35-year-old school enrol.	1969	13.59% (5.46%)	148	11.05% (4.20%)	1821	13.79% (5.50%)
19 to 21-year-old unemployment rate	1969	16.88% (10.26%)	148	15.53% (10.85%)	1821	16.99% (10.21%)
23 to 35-year-old unemployment rate	1969	8.13% (4.24%)	148	7.39% (3.64%)	1821	8.19% (4.29%)
high school dropout income (25-30 y.o.)	1965	\$9,702 (\$5,150)	148	\$9,961 (\$4,765)	1817	\$9,681 (\$5,180)
high school educated income (25-30 y.o.)	1969	\$14,428 (\$4,738)	148	\$14,610 (\$4,218)	1821	\$14,414 (\$4,778)
post-secondary educated income (25-30 y.o.)	1969	\$25,364 (\$6,338)	148	\$23,064 (\$5,102)	1821	\$25,550 (\$6,393)
19 to 21-year-olds moving in past year	1969	32.4% (14.8%)	148	35.5% (12.1%)	1821	32.1% (15.0%)
23 to 35-year-olds moving in past year	1969	29.1% (7.9%)	148	29.3% (6.5%)	1821	29.1% (8.1%)
value of oil reserves (reserves×price)	1969	\$1.106 bn (\$7.167 bn)	148	\$14.720 bn (\$22.045 bn)	1821	0 (0)

4 Results

4.1 Unemployment and Recessionary Effects

Figure 1 graphs the 19 to 21-year-old undergraduate enrolment rate alongside the unemployment rate for 19 to 21-year-olds in the sample. Interestingly, enrolment rates appear to have risen significantly over the period leading up to the 2008 Great Recession as the unemployment fell to a pre-recession low. The growth in the enrolment rate began to slow prior to the onset of the recession and does not appear to have been particularly affected by the subsequent rise in employment.

Figure 1: Undergraduate Enrolment Among 19 to 21-Year-Olds and the 19 to 21-Year-Old Unemployment Rate

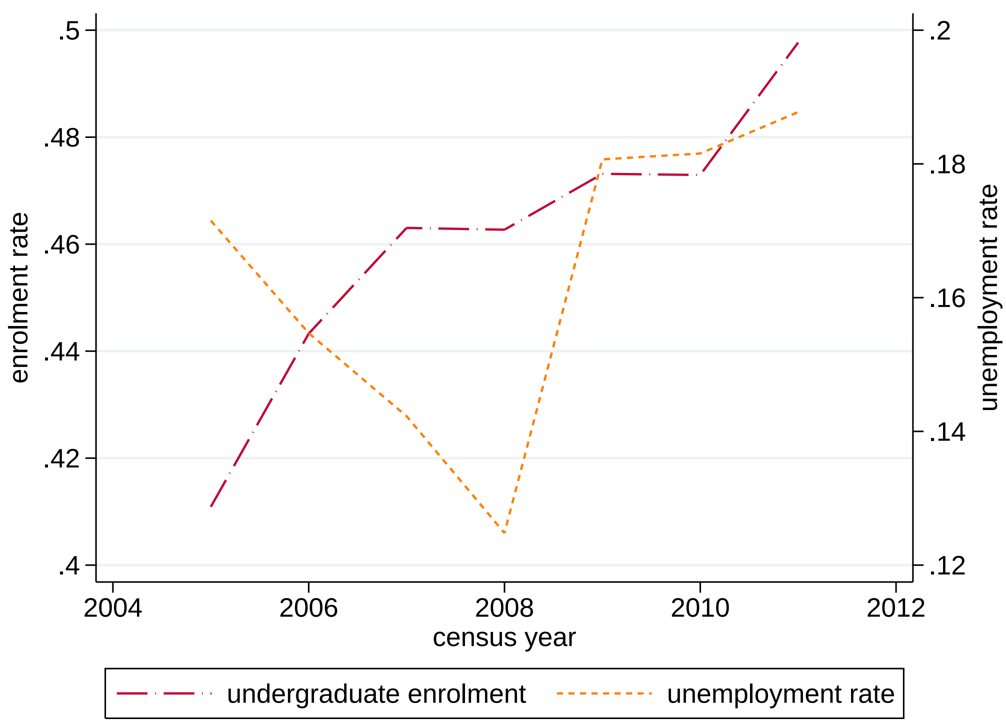

Figure 2 displays the enrolment rate among older workers aged 23 to 35 against the unemployment rate for the same age group. The graph shows a much tighter relationship between the unemployment and enrolment rates for this group. After

Figure 2: Enrolment Among 23 to 35-Year-Olds (All School Types) and the 23 to 35-Year-Old Unemployment Rate

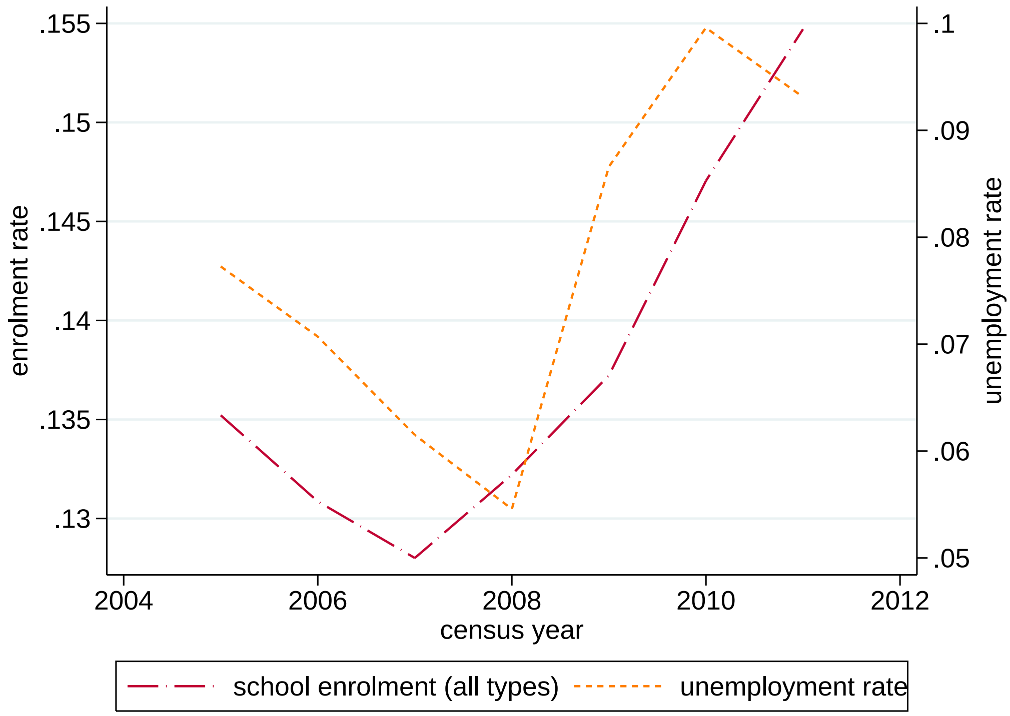

the onset of the recession the enrolment rate appears to have followed the rising unemployment rate with a one year lag. Prior to the recession, the unemployment rate and enrolment rate both decreased, however, the decline in enrolment appears to have preceded that of unemployment. Furthermore, the enrolment rate bottomed out prior to the onset of the recession. This suggests that the pre-recession decline in unemployment is unlikely to have been driven by changes in employment prospects.

Table 2 reports the coefficients and average partial effects (APE) from a set of probit regressions of enrolment status on local PUMA-level unemployment rates and their lagged values. Column (1) reports the results for college undergraduate enrolment among 19 to 21-year-olds as a function of the PUMA-level 19 to 21-year-old unemployment rate and its lag. Column (2) reports school enrolment of all types among 23 to 35-year-olds as a function of the PUMA-level 23 to 35-year-old unemployment rate and its lag. Since there is a risk that PUMAs with chronically low investment in human capital might be predisposed to having higher rates of unemployment, PUMA dummies were included in both regressions to control for pre-existing conditions. The APE estimates suggest that a one percentage point increase in the lagged unemployment rate among 19 to 21-year-olds corresponds with a 0.0685 percentage point increase in the probability of undergraduate enrolment, or a 0.15 percent increase in the enrolment rate. The estimated relationship for 23 to 35-year-olds is much more pronounced with a one percentage point increase in the unemployment rate or its lag corresponding with 0.104 or 0.0874 percentage point increase in the probability of school enrolment, respectively. These estimates represent 0.77% and 0.64% increases relative to the average enrolment rate for this age group. All of these estimates are statistically significant at the 99% level.

Columns (3) and (4) report the same probit regressions with the addition of year dummies and a proxy for migration into the PUMA which captures the proportion of people in the PUMA that moved in the past year. After the addition of these variables to the probit regressions, the coefficient on the lagged 19 to 21-year-old unemployment rate falls to just under half its previous value and ceases to be statistically significant. Of particular note is the fact that the coefficient for the contemporaneous unemployment rate becomes negative and statistically significant at the 99% level. The APE estimate suggests that a one percentage point increase in the unemployment rate corresponds with a 0.0583 percentage point decrease in undergraduate enrolment among 19 to 21-year-olds. This finding is consistent with the (albeit not statistically significant) findings of Card and Lemieux (2000) and could indicate that high unemployment or tighter credit conditions greatly hindered students' abilities to finance their education. For the 23 to 35-year-old age group, the addition of year dummies and migration controls renders the unemployment and lagged unemployment coefficients considerably smaller in magnitude and statistically insignificant. Since high

Table 2: Probit Regressions of Enrolment Probabilities on PUMA-level Unemployment Rates

	(1)		(2)		(3)		(4)	
Variables	Pr(*college*)	APE	Pr(*school*)	APE	Pr(*college*)	APE	Pr(*school*)	APE
(19-21 y.o. unempl. rate)$_t$	0.00168	0.000620	.	.	-0.158***	-0.0583***	.	.
	(0.0559)	(.020588)	.	.	(0.0590)	(.0217062)	.	.
(19-21 y.o. unempl. rate)$_{t-1}$	0.186***	0.0685***	.	.	0.0802	0.0295	.	.
	(0.0569)	(.0209643)	.	.	(0.0595)	(.021888)	.	.
(23-35 y.o. unempl. rate)$_t$	.	.	0.471***	0.104***	.	.	0.0121	0.00265
	.	.	(0.0839)	(0.0185)	.	.	(0.101)	(0.0222)
(23-35 y.o. unempl. rate)$_{t-1}$	.	.	0.397***	0.0874***	.	.	0.0602	0.0133
	.	.	(0.0861)	(0.0190)	.	.	(0.101)	(0.0223)
moving rate	.	.	.	.	-0.327	-0.120	0.0386	0.00850
	.	.	.	.	(0.253)	(0.0932)	(0.149)	(0.0328)
constant	-0.274***	.	-1.405***	.	-0.255***	.	-1.401***	.
	(0.0873)	.	(0.0550)	.	(0.0962)	.	(0.0599)	.
PUMA dummies	Yes		Yes		Yes		Yes	
Year dummies	No		No		Yes		Yes	
Observations	92,180		390,661		92,180		390,661	
Pseudo R^2	0.0697		0.0189		0.0709		0.0192	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

in-migration should correspond with a greater PUMA moving rate, if migratory responses to unemployment rates are introducing bias into the estimated unemployment-enrolment relationship then one would expect to find a negative relationship between the moving rate and enrolment. However, the moving rate coefficient is only negative in the 19 to 21-year-old enrolment regression and is not statistically significant for either age group. Given the low statistical significance on the moving rate coefficients and the high statistical significance on several of the year dummies (not reported), this change is likely being driven primarily by the inclusion of the year dummies[5].

The dramatic reduction in the magnitude of the estimated unemployment-enrolment relationship among 23 to 35-year-olds after the introduction of year dummies is surprising, especially given the trends visible in Figure 2 and earlier findings by Betts and McFarland (1995). It is possible that the unemployment-enrolment relationships estimated in columns (1) and (2) of Table 2 were driven by changes during the recession other than the unemployment rate specifically or that the severity of the recession resulted in a change in the unemployment-enrolment relationship (for example, as a result of the extended duration of high unemployment, changes in consumer confidence, or increased difficulty financing education). Both of these possibilities could explain why the year dummies better explain the variation in enrolment than the unemployment rate.

Table 3 presents probit regression results from an examination of the role of the Great Recession in explaining the observed variation in enrolment rates over the period. In each of the four probits estimated, the individual enrolment probability is regressed on the unemployment rate, a dummy set equal to 1 if the survey year was 2008 or later, an interaction term between these two variables, and a one year lag for each of the three preceding variables. Probit regression (1) focuses on undergraduate enrolment among 19 to 21-year-olds while regression (2) examines school enrolment

[5]Further unreported probit regressions confirm that exclusion of the moving rate dummy does not substantially change the coefficients on the unemployment rate variables.

among 23 to 35-year-olds. These regressions are repeated in columns (3) and (4) with the addition of year dummies and migration controls.

Probit regressions (1) and (3) both yield familiar negative estimates for the estimated unemployment-undergraduate enrolment relationship among 19 to 21-year-olds. The average partial effects estimates indicate that a one percentage point increase in the current unemployment rate corresponds with roughly a 0.07 percentage point decrease in enrolment which is statistically significant at the 95% level. There is also a positive relationship between the lagged interaction between the 2008 crisis dummy and the unemployment rate which is statistically significant at the 90% level. While this does suggest that after the onset of the recession high unemployment was more likely to encourage post-secondary enrolment (or, at least less likely to discourage enrolment), the magnitude of this estimate falls after the introduction of year dummies and migration controls. Probit regression (2) indicates that, with a confidence level of 95%, 23 to 35-year-olds were 0.068 percentage points more likely to be in school when the lagged unemployment rate increased by one percentage point in the lagged 2008 crisis period, however the estimated relationship falls in regression (4) with the introduction of control variables.

The fact that stands out the most in the Table 3 results is that the lagged 2008 crisis period indicator is statistically significant in all four probit regressions at the 99%, 90%, 99%, and 99% levels, respectively. The estimates suggest that the recession resulted in a roughly 2.0-5.1 percentage point increase increase in undergraduate enrolment among 19 to 21-year-olds and a 0.5-1.4 percentage point increase in school enrolment among 23 to 35-year-olds. This represents a roughly 4.6-11.8% increase in post-secondary enrolment relative to the mean among 19 to 21-year-olds, and a 3.7-10.3% increase relative to mean school enrolment for 23-35-year-olds. However, the upper end of the estimate for 19 to 21-year-olds should be interpreted with caution given the graphical evidence in Figure 1.

Table 3: Probit Regressions of Enrolment Probabilities on PUMA-level Unemployment Rates with 2008 Financial Crisis Interactions

Variables	(1) Pr(*college*)	(1) APE	(2) Pr(*school*)	(2) APE	(3) Pr(*college*)	(3) APE	(4) Pr(*school*)	(4) APE
[(19-21 y.o. unempl. rate)	0.0758	0.0279	.	.	0.0475	0.0175	.	.
$\times$*(crisis year)]*$_t$	(0.107)	(0.0394)	.	.	(0.107)	(0.0394)	.	.
[(19-21 y.o. unempl. rate)	0.180*	0.0662*	.	.	0.0654	0.0241	.	.
$\times$*(crisis year)]*$_{t-1}$	(0.0961)	(0.0354)	.	.	(0.0983)	(0.0362)	.	.
(19-21 y.o. unempl. rate)$_t$	-0.199**	-.0733**	.	.	-0.192**	-0.0706**	.	.
	(0.0961)	(0.0354)	.	.	(0.0966)	(0.0355)	.	.
(19-21 y.o. unempl. rate)$_{t-1}$	0.0497	0.0183	.	.	0.0459	0.0169	.	.
	(0.0775)	(0.0285)	.	.	(0.0778)	(0.0286)	.	.
[(23-35 y.o. unempl. rate)	.	.	-0.0733	-0.0161	.	.	-0.0186	-0.00409
$\times$*(crisis year)]*$_t$	.	.	(0.173)	(0.0381)	.	.	(0.174)	(0.0383)
[(23-35 y.o. unempl. rate)	.	.	0.311**	0.0684**	.	.	0.115	0.0252
$\times$*(crisis year)]*$_{t-1}$	.	.	(0.152)	(0.0334)	.	.	(0.159)	(0.0349)
(23-35 y.o. unempl. rate)$_t$	.	.	0.0953	0.0210	.	.	0.0188	0.00414
	.	.	(0.163)	(0.0358)	.	.	(0.165)	(0.0364)
(23-35 y.o. unempl. rate)$_{t-1}$	.	.	0.0757	0.0167	.	.	-0.00752	-0.00165
	.	.	(0.136)	(0.0298)	.	.	(0.138)	(0.0303)
(crisis year)$_t$	-0.0245	-0.00901	0.0135	0.00296	0.00347	0.00128	0.00749	0.00165
	(0.0204)	(0.00749)	(0.0133)	(0.00294)	(0.0223134)	(0.0082116)	(0.0147)	(0.00324)
(crisis year)$_{t-1}$	0.0545***	0.0201***	0.0233*	0.00513*	0.140***	0.0514***	0.0644***	0.0142***
	(0.0195)	(0.00719)	(0.0127)	(0.00279)	(0.0223)	(0.00820)	(0.0157)	(0.00345)
moving rate	.	.	.	.	-0.319	-0.117	0.0419	0.00923
	.	.	.	.	(0.254)	(0.0933)	(0.149)	(0.0329)
constant	-0.274***	.	-1.398***	.	-0.248***	.	-1.398***	.
	(0.0881)	.	(0.0559)	.	(0.0967)	.	(0.0607)	.
PUMA dummies	Yes		Yes		Yes		Yes	
Year dummies	No		No		Yes		Yes	
Observations	92,180		390,661		92,180		390,661	
Pseudo R^2	0.0703		0.0191		0.0709		0.0192	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

4.2 Variation in Income Premia by Local Oil Intensity

In order to determine the suitability of the oil intensity instrument for estimating the effect of wage premia on enrolment decisions, it must first be established whether or not local oil intensity does, in fact, affect earnings and earnings differentials by education level. Figure 3 shows average wage/salary income among 25 to 30-year-olds by educational attainment in oil-rich PUMAs plotted alongside the oil price over the 2005-2011 period. It was decided to focus on incomes for 25 to 30-year-olds since this group is considerably more likely to be engaged in work as opposed to attending school. Furthermore, the earnings among this group are likely to be more representative of the post-schooling prospects of youths than those of older cohorts with more work experience and at greater risk of generational differences.

Figure 3: Mean Wage/Salary Income by Education Level for 25 to 30-Year-Olds in Oil-Rich PUMAs

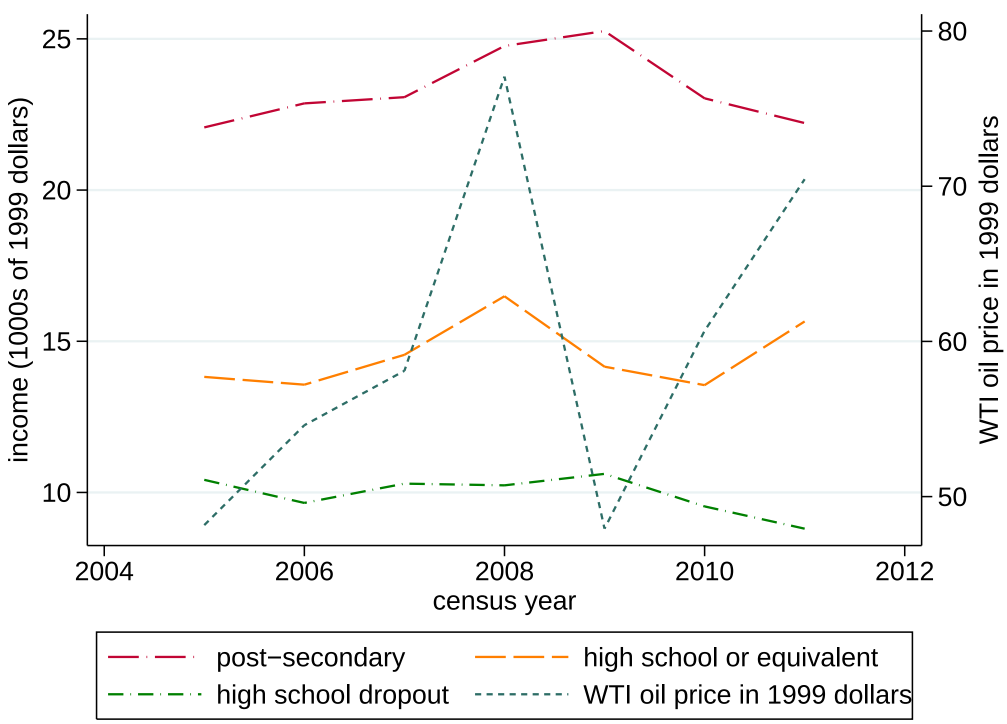

The incomes of both those completing high school and post-secondary education rose by several thousand dollars over the course of the 2006-2008 rise in oil prices. Average earnings among high school graduates then began to decrease after oil prices fell from the 2008 peak while earnings among those completing post-secondary education took an additional year before beginning to decline. With the second major increase in oil prices after 2009, earnings among high school graduates began to climb again, although there was no similar recovery for post-secondary graduates. Earnings among 25 to 30-year-old high school drop-outs remained stable over the 2005-2011 period.

Figure 4 provides a similar graph for PUMAs without significant oil reserves. There appears to be no variation in the average earnings of 25 to 30-year-olds with various levels education that coincides with changes in the oil price. In fact, average

Figure 4: Mean Wage/Salary Income by Education Level for 25 to 30-Year-Olds in PUMAs Without Significant Oil Reserves

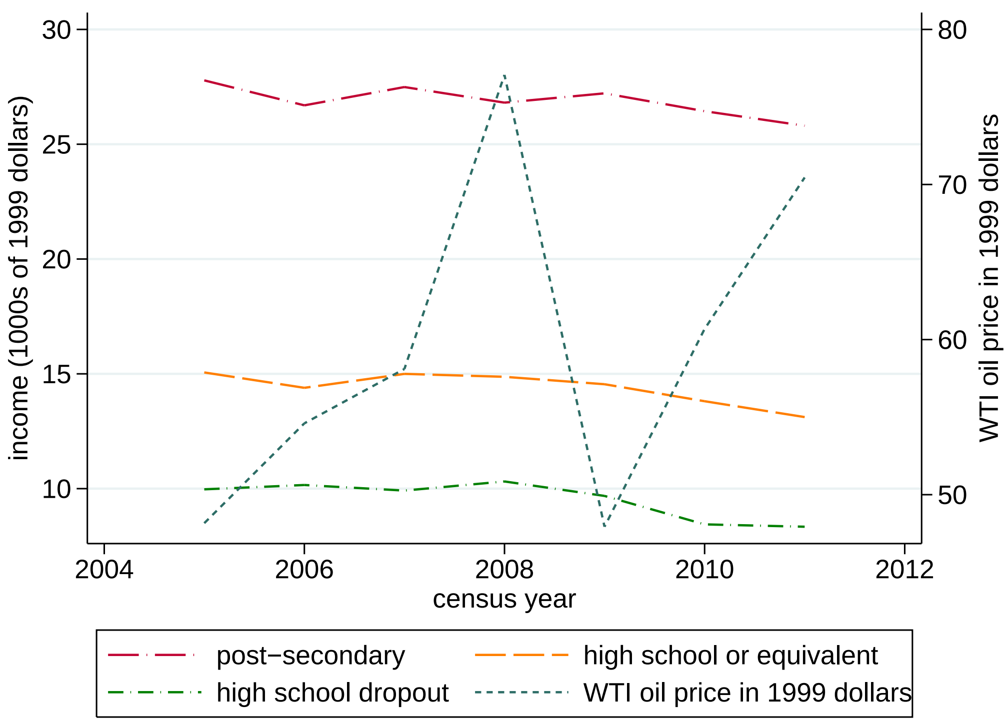

earnings are essentially flat for all three educational attainment groups over the 2005-2011 period, which is consistent with the predictions made in this paper.

While the contrast between Figures 3 and 4 suggests that earnings vary in response to oil prices in regions with substantial oil reserves, ultimately it is variation in earnings premia, and not absolute earnings, that is relevant for the present analysis. Figure 5 reveals that exposure to fluctuations in oil prices in regions with significant reserves can have substantial effects on wage premia. The graph plots average earnings premia for 25 to 30-year-old high school graduates relative to high school drop-outs for both oil-rich and oil-poor PUMAs. There appears to be a very strong relationship between the high school earnings premium and the oil price in the oil-rich PUMAs which is not mirrored in PUMAs lacking in oil reserves. While this paper does not

Figure 5: Income Premium for 25 to 30-Year-Olds with a High School Education Relative to High School Drop-outs (for PUMAs with and without Significant Oil Reserves)

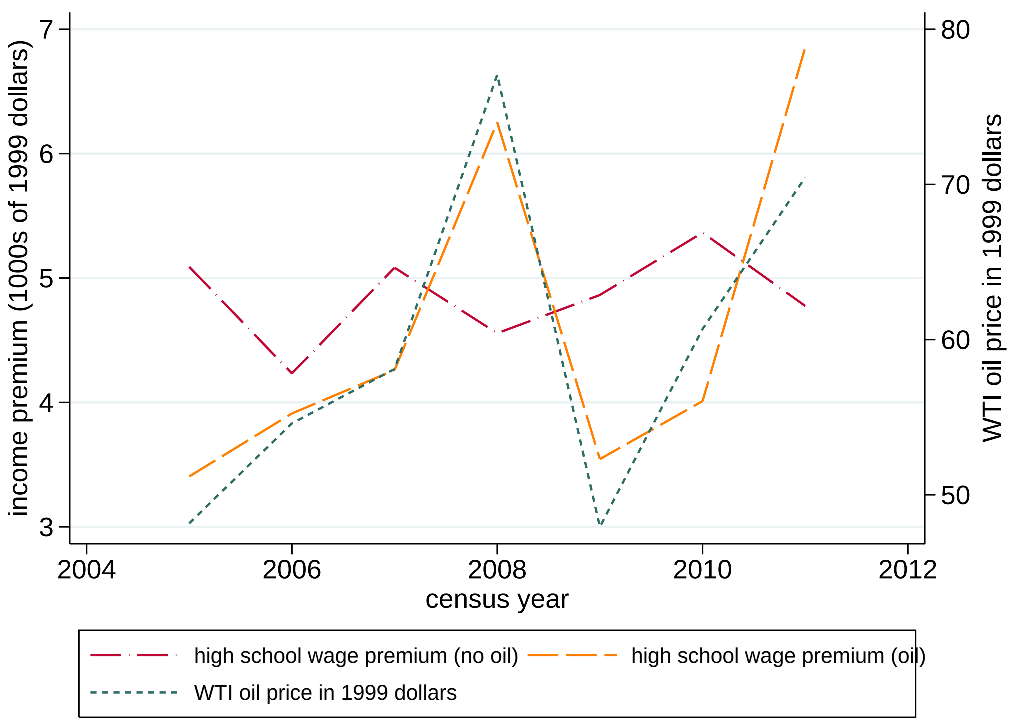

focus on high school drop-out decisions, the graph is, nonetheless, illustrative of the rationale for using the oil intensity instrument as a source of plausibly exogenous variation in earnings premia.

Figure 6 graphs the post-secondary earnings premia (relative to high school graduate earnings) for 25 to 30-year-olds in both PUMAs with and without significant oil reserves alongside oil prices. The earnings premium remains fairly flat in oil-poor PUMAs but appears to be inversely related to the oil price in the oil-rich PUMAs. Interpreted in conjunction with the previous figure, this suggests that any wage gains associated with increased oil prices were larger for high school graduates than their more educated counterparts. These findings are consistent with predictions made earlier in this paper and the findings of Black et al. (2005) for coal-rich counties in the Appalachian states.

Figure 6: Income Premium for 25 to 30-Year-Olds with a Post-Secondary Education Relative to High School Graduates (for PUMAs with and without Significant Oil Reserves)

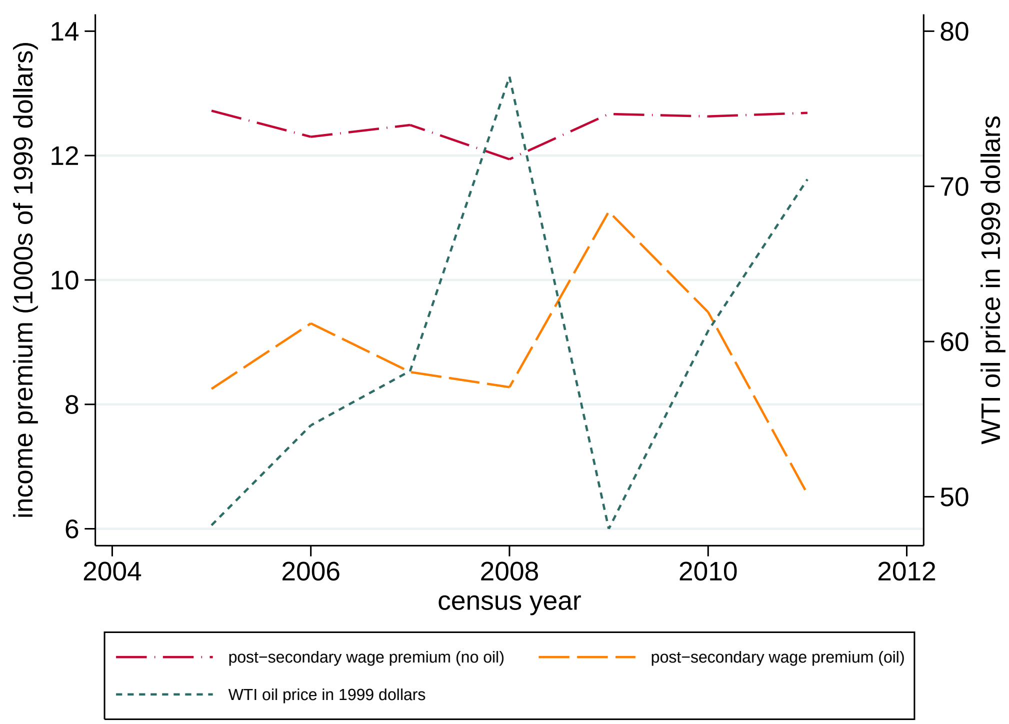

Table 4 reports the results for the regression outlined in equation (2). Individual wage/salary income was regressed on dummy variables specifying each individual's highest level of completed schooling, the value of oil reserves instrument, and interaction terms between these dummies and the instrument. The value of oil reserves instrument was normalized by dividing it by \$14.720 bn—the mean value among PUMAs with significant oil reserves. The omitted educational attainment group is high school drop-outs. The results in column (1) indicate that high school graduates, individuals with some post-secondary education, and post-secondary graduates between the ages of 25 and 30 earned roughly \$4,968, \$7,974, and \$18,220 more than their high school dropout counterparts, respectively. These amounts increased to \$6,030, \$9,346, and \$20,810 after controlling for potential work experience in column (2) and remained essentially unchanged after PUMA and year dummies were added in column (3). These estimates are all highly significant at the 99% level.

The coefficients on the instrument are negative but insignificant for the first two regressions. After introducing PUMA and year dummies the coefficient becomes positive and statistically significant at the 99% level—likely due to the fact that PUMA-level controls address the fact that oil-rich regions might differ systematically from low-oil regions in terms of factors such as population density. Given the normalization applied to the instrument, the coefficient of 1.414 suggests that, after controlling for education, potential experience, and PUMA and year effects, a state with the average value of oil reserves among the oil-rich PUMAs would be expected to have average incomes \$1,414 greater than a similar PUMA without significant reserves.

Perhaps the most interesting features of Table 4 are the estimated coefficients for the interaction terms between the instrument and the educational attainment dummies. The coefficients on the interaction between the value of oil reserves instrument and the some-college-attainment dummy in all three regressions indicate that the income premium for individuals with some post-secondary education relative to

Table 4: Regressions of Wage/Salary Income on Education Level and Normalized Oil Intensity (1000s of 1999 dollars)

Variables	(1) *income*	(2) *income*	(3) *income*
(value of oil reserves) × (max. educ. = high school)	0.0252	0.0225	0.0329
	(0.185)	(0.186)	(0.185)
(value of oil reserves) × (max. educ. = some college)	-0.520***	-0.532***	-0.523***
	(0.186)	(0.187)	(0.188)
(value of oil reserves) × (min. educ. = college)	-1.376***	-1.430***	-1.228***
	(0.236)	(0.237)	(0.233)
(max. educ. = high school)	4.968***	6.030***	5.902***
	(0.101)	(0.115)	(0.115)
(max. educ. = some college)	7.974***	9.346***	8.879***
	(0.106)	(0.130)	(0.130)
(min. educ. = college)	18.22***	20.81***	19.48***
	(0.115)	(0.178)	(0.176)
value of oil reserves	-0.0130	-0.00711	1.414***
	(0.125)	(0.126)	(0.480)
potential experience		0.418***	0.427***
		(0.0203)	(0.0202)
constant	9.100***	4.078***	7.534***
	(0.0746)	(0.256)	(0.821)
PUMA dummies	No	No	Yes
Year dummies	No	No	Yes
Observations	208,655	208,655	208,655
R-squared	0.116	0.117	0.142

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

high school drop-outs is estimated to be about $500 dollars smaller in an average oil-rich PUMA relative to an oil-poor one. This dynamic is even more pronounced for post-secondary graduates who were estimated to make, on average, $1,228 less in an average oil-rich PUMA than in an oil-poor one after controlling for potential experience, and PUMA and year effects. Furthermore, it is worth noting that the addition of year and PUMA dummies had a minimal effect on the coefficients of the interaction terms, suggesting that pre-existing PUMA conditions and year-specific shocks other than oil price fluctuations explain very little of the variation in earnings premia between oil-rich and oil-poor regions. These findings suggest that the value of oil reserves at the PUMA-level could provide a suitable source of variation in earnings premia in the IV regressions to follow.

4.3 The Effects of Income Premia on Enrolment Decisions

Figure 7 shows the undergraduate enrolment rates among 19 to 21-year-olds in PUMAs with and without significant oil reserves alongside oil prices. The enrolment rates in these two groups appear to track each other prior to the Great Recession—a period in which oil prices were rising. After peaking in 2007, the enrolment rate in the oil-rich regions began to decline relative to the low-reserve regions as the oil price peaked in 2008. This is consistent with the declining post-secondary earnings premium that was found to accompany rising oil intensity. Enrolment in oil-rich regions began to climb again after the oil price bottomed out in 2009. However, it should be noted that the relationship between the oil price and the enrolment rate in oil-rich PUMAs is not overwhelmingly strong. The recession's potential impact on the ability of individuals to finance further education cannot be ruled at as a potential explanation for these changes. Nonetheless, it is not obvious what other factors could be driving the post-2007 divergence in enrolment rates between high- and low-oil reserve areas.

As Betts and McFarland (1995) point out, 23 to 35-year-olds are likely to face very different enrolment incentives than their younger counterparts. Older workers will have a lower net present value of future earnings, will likely face greater forgone earnings upon returning to school, and are likely to have a higher discount rate given that they already chose to discontinue their schooling in the past. It follows from this that their enrolment decisions should be less sensitive to changes in earnings premia. As such, this group can serve as a sort of placebo. A graph of the 23-35-year-old enrolment rate by oil-rich and oil-scarce PUMAs alongside the oil price can be found in the appendix. As expected, the graph does not reveal any notable relationship between the oil price and enrolment in oil-rich PUMAs.

Figure 7: Undergraduate Enrolment Rates Among 19 to 21-Year-Olds in Oil-Rich and Oil-Poor PUMAs

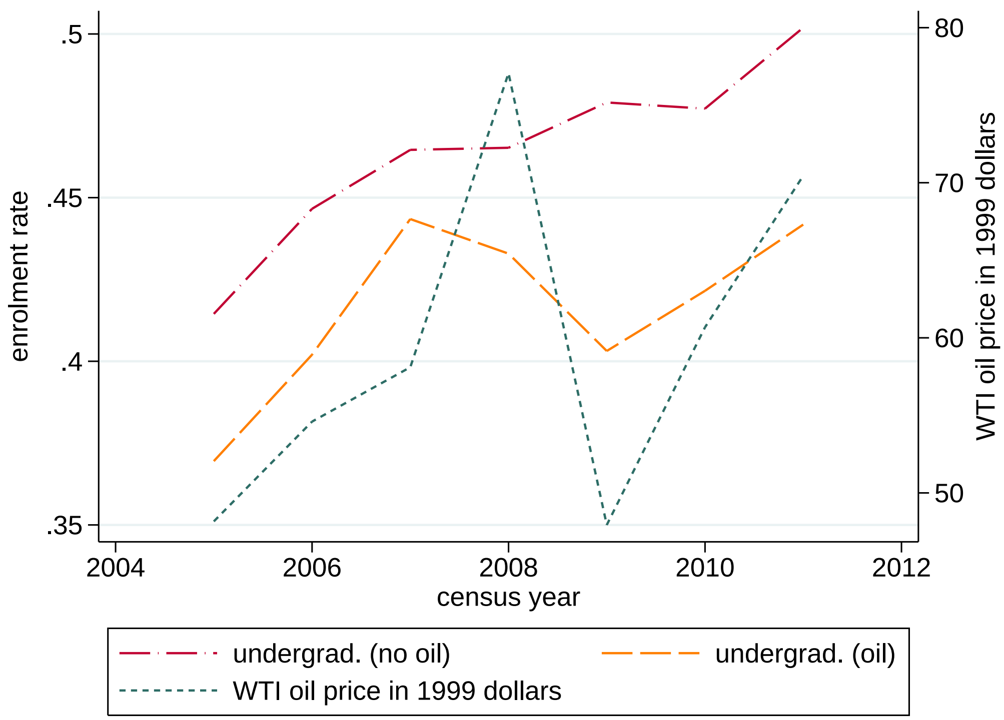

Table 5 shows the results for OLS and IV regressions of the 19 to 21-year-old undergraduate enrolment rate on the post-secondary income premium computed for 25 to 30-year-olds. Column (1) reports the OLS regression results with PUMA fixed effects and year dummies. The estimated coefficient on the post-secondary income premium variable is extremely small at 0.000439—which indicates that a $1,000 increase in the income premium corresponds with a 0.04 percentage point increase in enrolment—and is not statistically significant. Columns (2) and (3) report the IV estimates of the enrolment-income premium relationship with the income premium instrumented by the normalized PUMA value of oil reserves variable and its lag. Column (3) adds a proxy for migration—the moving rate—to the IV regression in column (2) which captures the average number of individuals in the PUMA who moved in the past year. However, none of the post-secondary earnings premium coefficients are statistically significant. In fact, the absolute values of their z-scores do not exceed 0.5 in either IV regression. Nonetheless, it is worth noting the the magnitude of the point estimate on the post-secondary earnings premium variable increases a hundred-fold to about 0.04, indicating that a $1,000 increase in the post-secondary earnings premium corresponds with a very substantial 4 percentage point increase in the post-secondary enrolment rate. However, the IV estimates are far too imprecise to arrive at any reliable conclusions regarding the enrolment-earnings premium relationship or to determine whether or not the OLS estimates are biased.

The poor performance of the panel IV regressions is almost certainly due to the weakness of the oil intensity and lagged oil intensity instruments since the magnitude of the point estimates is non-trivial. The results of the first stage estimates are also reported in Table 5 and reveal that the coefficients on the instruments are imprecisely estimated. The results of further tests (not reported) indicate that the correlations between the income premium and the oil intensity variable and its lag are -0.0529 and -0.0533, respectively, and OLS regression of the earnings premium on these two

Table 5: Panel IV Regressions of PUMA-level Enrolment Rates on Earnings Premia Instrumented by Normalized Values of PUMA Oil Reserves

Variables	OLS (1) *enrolment rate*	IV (2) *enrolment rate*	IV (3) *enrolment rate*
earnings premium	0.000439	0.0423	0.0423
	(0.000362)	(0.163)	(0.163)
moving rate			-0.0125
			(0.128)
constant	0.388***	-0.0553	-0.0505
	(0.00698)	(1.824)	(1.791)
		First Stage	
		(2) *earnings premium*	(3) *earnings premium*
(value of reserves)$_t$		0.274	0.296
		(2.059)	(2.060)
(value of reserves)$_{t-1}$		-0.396	-0.378
		(2.034)	(2.036)
moving rate			0.617
			(1.800)
constant		11.189***	10.970***
		(0.437)	(0.774)
Year Dummies	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes
Observations	1,969	1,685	1,685
Number of id	284	281	281
R^2	0.0183	0.0002	0.0002

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

instruments yields an F-score of $F(2, 1682) = 2.42$—below the rule-of-thumb of 10. It can only be concluded that these panel IV regression results are insufficient to determine the relationship between earnings premia and post-secondary enrolment.

Despite the inconclusive IV results for the enrolment-earnings premium relationship, it is nonetheless interesting to note that the addition of the migration rate proxy to the regression in column (3) had no impact on the point estimate of the earnings premium coefficient. Overall, the panel IV regression results suggest that the relationship between post-secondary enrolment and migration rates is fairly weak given the statistically insignificant coefficient on the moving rate variable.

Conclusion

The goal of this study was to identify the ways in which labour market conditions impact school enrolment decisions among two age groups in the United States: 19 to 21-year-olds and 23 to 35-year-olds. In order to quasi-experimentally determine the effects of earnings premia on enrolment decisions, it was decided to focus the analysis on six U.S. states—Texas, New Mexico, Oklahoma, Louisiana, Mississippi, and Colorado—that were deemed suitable because of their oil reserves and PUMA sizes. This paper has endeavoured to determine the causal effect of unemployment rates on individual enrolment decisions using the plausible exogeneity of the unemployment rate, and the causal effect of the post-secondary earnings premium on post-secondary enrolment through the use of plausibly exogenous variation in the value of local oil reserves in an instrumental variables approach.

This study has found evidence that enrolment rates among 23 to 35-year-olds rose in the 2009-2011 period, after the onset of the recent recession, which is consistent with earlier findings by Betts and McFarland (1995). Some evidence was found that post-secondary enrolment among 19-21-year-olds also rose after the onset of the recession, however the enrolment rate was also found to be negatively related to the unemployment rate for most of the 2005-2011 period, perhaps as a result of greater difficulties individuals may have experienced in financing their education.

As part of the process of instrumenting earnings premia by education using an oil shock instrument, it was necessary to examine the effects of exposure to oil prices on these premia. This study has found a significant relationship between the value of local oil reserves (the product of local oil reserves and the oil price) and local wage premia. The wage premia for individuals with some post-secondary education and for post-secondary graduates were both found to be significantly smaller in PUMAs with a high value of oil reserves. Overall, these findings support the conjecture that oil shocks affect earnings differentials by education in exposed regions. These results are consistent with those of Acemoglu et al. (2013) who find positive effects on incomes in oil-rich regions, and Black et al. (2005) who find effects on earnings differentials by education in coal-rich counties, albeit among different education levels.

Despite the findings of economically and statistically significant variation in educational earnings premia by local oil intensity, the results of the IV regressions presented in this paper were inconclusive. The precision of the estimates was very low—likely as a result of weak instruments, given the much larger magnitude of the IV estimates in comparison to the OLS estimates. In future research, the problem with weak oil intensity instruments could be addressed by using the non-public version of the American Community Survey or alternative datasets that provide more localized geographic units, such as counties, rather than Public Use Microdata Areas (PUMAs). This would represent an improvement over the methodology used in this paper which yielded 21 oil-rich PUMAs relative to a total of 281 PUMAs appearing across the 2005-2011 period in the sample. However, a sample of the ACS much larger than 1% should be used when increasing the geographic precision beyond the PUMA level.

Future extensions to this research could involve refining the oil intensity instrumental variable approach, as outlined above, in order to arrive at more reliable results pertaining to the enrolment-earnings premium relationship. Alternatively, improved results could be obtained by exploiting alternative instruments such as the Mexican export demand shocks used by Atkin (2012). Future research could also explore the potential impact of labour market conditions on more subtle education decisions such as college major choices, and whether students delay their graduation in periods of high unemployment, among other topics.

References

Acemoglu, Daron, Amy Finkelstein, and Matthew Notowidigdo (2013), "Income and health spending: Evidence from oil price shocks." *Review of Economics and Statistics*, 95(4), 1079–1095.

Atkin, David (2012), "Endogenous skill acquisition and export manufacturing in mexico." *NBER Working Paper 18266*, http://www.nber.org/papers/w18266.

Becker, Gary (1962), "Investment in human capital: A theoretical analysis." *Journal of Political Economy*, 70(5), 9–49.

Betts, Julian and Laurel McFarland (1995), "Safe port in a storm: The impact of labor market conditions on community college enrollments." *Journal of Human Resources*, 30(4), 741–765.

Black, Dan A., Kermit Daniel, and Seth G. Sanders (2002), "The impact of economic conditions on participation in disability programs: Evidence from the coal boom and bust." *American Economic Review*, 92(1), 27–50.

Black, Dan A., Terra G. McKinnish, and Seth G. Sanders (2005), "Tight labor markets and the demand for education: Evidence from the coal boom and bust." *Industrial and Labor Relations Review*, 59(1), 3–16.

Card, David and Thomas Lemieux (2000), "Dropout and enrollment trends in the post-war period: What went wrong in the 1970s?" *NBER Working Paper No. 7658*, http://www.nber.org/papers/w7658.pdf.

Chaplin, Duncan, Mark Turner, and Andreas Pape (2003), "Minimum wages and school enrollment of teenagers: a look at the 1990's." *Economics of Education Review*, 22, 11–21.

Clark, Damon (2011), "Do recessions keep students in school? the impact of youth unemployment on enrolment in post-compulsory education in england." *Economica*, 78, 523–545.

Dellas, Harris and Vally Koubi (2003), "Business cycles and schooling." *European Journal of Political Economy*, 19, 843–859.

Dellas, Harris and Plutarchos Sakellaris (2003), "On the cyclicality of schooling: Theory and evidence." *Oxford Economic Papers*, 55(1), 148–172.

Edwards, Linda (1976), "School retention of teenagers over the business cycle." *Journal of Human Resources*, 11(2), 200–208.

Emery, Herb, Ana Ferrer, and David Green (2011), "Long term consequences of natural resource booms for human capital accumulation." *Canadian Labour Market and Skills Research Network Working Paper No. 74*, http://www.clsrn.econ.ubc.ca/workingpapers/CLSRN Working Paper no. 74 – Emery, Ferrer, Green.pdf.

Guo, Naijia (2013), "The impact of an early career recession on schooling and lifetime welfare." *University of Pennsylvania Economics Job Market Paper*.

Gustman, Alan and Thomas Steinmeier (1981), "The impact of wages and unemployment on youth enrollment and labor supply." *Review of Economics and Statistics*, 63(4), 553–560.

Light, Audrey (1996), "Hazard model estimates of the decision to reenroll in school." *Labour Economics*, 2(4), 381–406.

Loken, Katrine V. (2007), "Family income and children's education: Using the norwegian oil boom as a natural experiment." *University of Bergin Working Paper*, http://www.uib.no/filearchive/No. 03–07.pdf.

Marchand, Joseph (2012), "Local labor market impacts of energy boom-bust-boom in western canada." *Journal of Urban Economics*, 71, 165–174.

Neumark, David and William Wascher (1995), "Minimum-wage effects on school and work transitions of teenagers." *American Economic Review*, 85(2), 244–249.

Steven Ruggles, J., Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek (2010), *Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]*. University of Minnesota, Minneapolis, MN.

van den Berg, Gerard J., Maarten Lindeboom, and France Portrait (2006), "Economic conditions early in life and individual mortality." *American Economic Review*, 96(1), 290–302.

Appendix

Figure 8: Enrolment Rate Among 23 to 35-Year-Olds (All School Types) for PUMAs with and without Significant Oil Reserves