Issue 10 Summer 2010

CMPO

Research in PUBLIC POLICY

Bulletin of the Centre for Market and Public Organisation

EDUCATION

Research evidence on 'free schools', closures and parental choice

ALSO IN THIS ISSUE:

Early life health · Measuring diversity

Research in PUBLIC POLICY

Bulletin of the Centre for Market and Public Organisation

Summer 2010

Contents

To listen to an audio interview with the author, visit: http://www.bristol.ac.uk/cmpo/audio/

Research evidence on school choice

Education remains one of the key policy areas for politicians and featured prominently in all three of the main parties' election manifestos. The focus of central government on raising standards is not going to go away, but the pressure to find ways of raising standards in the face of at best static budgets is likely to intensify over the short to medium term.

The three articles featured in this issue of *Research in Public Policy* are therefore particularly timely, as they cover a range of issues relevant to the future of the system of school choice in England.

Recent research on the education market has focused on two key aspects of competition: entry into and exit from the market; and the appropriate policy environment. The articles by Rebecca Allen and Bruce Sacerdote draw on international experience to provide evidence, respectively, on the potential effects of new schools entering the market and of failing schools closing.

Rebecca Allen's review of the evidence on the Swedish 'free school' reforms concludes that increasing entry in this way did have a positive effect on education outcomes. The benefits, however, are small; they are predominantly focused on children from highly educated families and they do not persist: scores are no higher in the end-of-school exams.

Bruce Sacerdote's results show that moving children out of failing schools can have a positive effect on their outcomes. This is particularly true for children with low prior attainment and it does not come at the cost of harming the performance of students in the schools that absorbed the new intakes.

These two articles provide evidence on the potential positive effects of exit and entry on educational outcomes. But the potential for any school choice system to improve academic standards relies on parents choosing schools on that basis. Parental preferences for different characteristics of schools fundamentally affect the outcome of any choice-based assignment mechanism.

The article by Deborah Wilson provides evidence on this issue. The research she reports on finds that parents do value academic standards, and that they also value particular peer groups, which creates pressure for socio-economic 'sorting'. A key finding, however, is that it is differences in location – and thereby the types of schools accessible to families – and not differences in preferences that drive differential pupil assignment across schools.

So parents do value academic standards, which suggests that a well-functioning school choice system has the potential to improve them. Entry into and exit from the education market are essential for such a system to function effectively. But a more fundamental problem may be addressing the inequality of access currently present in England's system of school choice.

Replicating Swedish 'free school' reforms in England

Rebecca Allen **of the Institute of Education summarises the research evidence on the impact of Sweden's educational reforms on pupils' academic achievement.**

The Conservatives' proposal to replicate Sweden's 'free school' reforms would be the most radical reform of education in England since the dismantling of selective schooling four decades ago. But the work of Swedish economists used to support the argument that the introduction of choice and competition has improved academic performance is less unambiguous than the Conservatives claim.

There are a number of related reforms to the supply side that have to be understood as background to the changes and which raise questions about the potential of free schools in the UK, particularly if pursued in isolation. Here I summarise evidence on the impact of the Swedish reforms on test scores, evaluating the relative merits of the research papers and explaining why they disagree in their findings.

The background to Sweden's reforms

In 1992, Sweden introduced a voucher scheme in which privately run (including for-profit) schools could receive public funding for each pupil they educated on the same terms as municipality schools. Like all market reforms of public services, the exact nature of the institutional structures, financing and regulation are critical to ensuring success.

But the important context for the changes in Sweden was a backdrop of radical supply-side reforms, which were intended to facilitate innovation and more efficient resourcing decisions. These included deregulating teachers' pay and conditions, decentralising school financing and increasing schools' discretion over curriculum, goal-setting and test regimes (see Björklund et al, 2005, for an overview of the reforms).

Each reform was almost the exact opposite of the New Labour education reforms that were taking place at the same time in the UK.

Today, about 10% of lower secondary aged pupils in Sweden choose to attend the privately run free schools, with places strictly allocated on a 'first-come, first-served' basis. First-come, first-served sounds like a sensible and fair admissions rule, but it can lead to parents queuing on streets overnight to ensure they are first in line on the day that the admissions process opens.

The best evidence on Sweden's reforms indicates small improvements in academic achievement in areas with more free schools

There has been large regional variation in the expansion of free schools. More have been established in urban, affluent and gentrifying areas and in those places with second-generation immigrant communities. Within these municipalities, more educated parents and second-generation immigrants are most likely to use the free schools, so the overall system is stratifying a little (Böhlmark and Lindahl, 2007). The largest group of free schools are for-profit providers of a general education, but special pedagogy, religious and special language/ethnic group schools are also prevalent.

Evidence of the impact on academic achievement

The Swedish schooling reforms are not straightforward to analyse because free schools were not set-up at random across the 290 municipalities: they are more prevalent where the municipality is politically supportive and offers high per pupil funding. Movement of pupils across municipality boundaries is also permissible, although not particularly common.

Evaluation of the reforms is especially difficult because Sweden does not routinely collect administrative test score and demographic data on all pupils in the country, as England does.

Externally marked test score data in maths, English and Swedish are available for around 30 municipalities, but for the remainder, researchers are restricted to using grade levels that are not consistently standardised across the country.

The research papers described below adopt a variety of strategies for dealing with the non-random assignment of the policy and the fact that areas with many free schools have demographic characteristics that make them systematically different. Two studies rely on cross-sectional data, hoping that the good quality of control variables (and the use of an instrument to predict free school supply) avoid confounding influences.

The most recent two studies use a municipality panel to measure the extent to which changes in the share of free schools in the municipality are associated with changes in pupil test scores. This latter approach requires less arduous identification assumptions since it controls for time-invariant demographic factors in the municipality.

There is still a problem of establishing causality as trends in social demographics (such as an influx of immigrants) are associated with trends in free school growth. But this can be dealt with by collecting data on time-varying demographic characteristics and/or by accounting for pre-reform trends in test scores.

Böhlmark and Lindahl

The most recent papers to evaluate the reforms (Böhlmark and Lindahl, 2007, 2008) are described first since they use the most robust data and methods, requiring relatively few identification assumptions. By describing these first, the relative drawbacks of the other papers' approaches can be understood.

The biggest advantages of this research come from the construction of a long panel of data from 1988/89 to 2002/03. This gives three years of pre-intervention trend and over a decade of post-reform data to apply a difference-in-differences approach with municipality fixed effects to compare changes in areas with large growth in free schools with changes in areas with smaller growth in free schools.

The sample of data is also very large: a 20% sample from the population of pupils across all municipalities. The quality of the background control variables is good, including parental education, income, age and immigrant status, although they do not include any measures of the child's prior attainment.

The researchers do rely on non-standardised attainment data, but they have externally marked test scores for a sample of pupils that are sufficient to confirm that biases in teacher assessment are not correlated with the policy reform.

Small positive effects are evident in both the private schools and – through competition – in municipality schools

The results show a moderately positive impact of free school growth on municipality academic performance at the end of ninth grade (the end of lower secondary school, when pupils are aged 15-16). This finding is convincing because it is consistently estimated across almost all subjects and model specifications. The biggest beneficiaries are children from highly educated families; the impact on low educated families and immigrants is close to zero.

By tracking siblings within families who differed in whether they attended municipality or free schools, the researchers show that the superior performance of areas with private schools is due both to the greater effectiveness of private schools and to municipality schools making improvements in response to school competition, with the latter likely to be more important than the former.

But the researchers also find that the advantages that children educated in areas with free schools have by age 16 do not translate into greater educational success in later life. Although there is some (weak) evidence that pupils in areas with many free schools are more likely to take an academic track in high school, they score no better in high school exit tests at the age of 18/19. They are also no more likely to participate in higher education than those who were schooled in areas without free schools.

The researchers explore a variety of explanations for this, but conclude that the educational advantages of school competition are simply too small to persist into any long-term gains for young people.

Björklund et al

The short panel of data for municipalities between 1998 and 2001 that Björklund et al (2004, 2005) use is only able to analyse the relationship between growth in private school share in a municipality and changes in test scores over a short period of time, with correspondingly less variation in the parameter of interest (just a one percentage point change in free schooling share between these dates) and no pre-reform data to account for pre-existing social trends.

The researchers compare estimates between a sample of around 30 municipalities for which they have good quality data and all 290 municipalities where data quality is poor. Overall, they do not find a consistently positive impact of free school share on educational attainment: they identify a small positive impact on English and Swedish attainment, but a zero or even negative impact in maths.

Their findings are not consistent across the sample and the population of municipalities, suggesting that there may be selection problems in the municipality sample. This is a significant observation about data quality since the following two studies both rely on this sample of 30 municipalities.

Ahlin

Ahlin (2003) estimates the impact of the share of private schooling on ninth grade test scores in a cross-section of 34 municipalities from 1997/8, hoping that the quality of her control variables are sufficient to avoid any confounding influences. This is the only study that includes the prior attainment of the pupil in sixth grade and further background controls, thus accounting for systematic differences in the levels of attainment across municipalities but not dealing with differences in expected rates of progress from sixth grade to ninth grade that are due to home background factors.

The educational advantages of school competition are too small to persist into any long-term gains for young people

Her findings reverse those of Björklund et al with quite large positive effects of private schools on overall municipality achievement in maths, but not in Swedish or English.

Sandström and Bergström

Sandström and Bergström (2005) were the first researchers to explore the impact of the free school reforms on overall academic standards in Sweden. Their finding of large positive gains to the reforms have been widely reported, and are surprising given that their data come from quite early in the reform period (1997/8) before growth in free schools became substantial.

The study relies on the largest number of identifying assumptions since they use a cross-section of only municipality schools in just 30 municipalities, using a parametric sample-selection correction to address composition changes caused by lack of data on pupils in free schools. They use a two-stage approach, with an instrument of political control predicting the municipality share of free schools.

Critics argue that the instrument may not meet the excludability criterion of predicting the growth of free schools but not directly determining education attainment because they are not able to control for most of the social factors in the municipality that explain household educational practices. Given this, it is hard to argue that their large positive finding should contribute to our current knowledge of the impact of the reforms.

Concluding remarks

The experience of Sweden is helpful, but necessarily limited, in the extent to which it can help predict the impact of school reforms in England. One reason for this is that the schools also underwent a radical decentralisation of the education system, which would seem to be critical for promoting diversity and productivity gains through experimentation in free schools.

Sweden also has fewer reasons to be concerned that a free school system will produce greater school stratification since the country's lower levels of income and skill inequalities mean there is far less need for parents to choose schools based on social composition. It is also possible that Sweden's stronger tradition of non-standard schooling (such as Steiner and Montessori schools) is leading to a greater diversity of provision than parents in England would ever demand.

Sweden's experience is limited in the extent to which it can predict the impact of comparable school reforms in England

The econometric evidence on the impact of the reforms suggests that, so far, Swedish pupils do not appear to be harmed by the competition from private schools, but the new schools have not yet transformed educational attainment in Sweden. Bunar (2009) argues that the growth in free schools in the first decade was too slow to bring a great transformation due to unclear regulations and uncertainty as to whether the ruling Social Democratic Party would further weaken the financial conditions for free schools.

In addition, the rising pupil population in Sweden during the 1990s meant that existing state schools did not lose pupils in great numbers as free schools opened and poorly performing schools did not need to close. In the past few years, overall demand for school places has fallen as the pupil population shrinks and the supply of free schools places has rapidly grown. So the prospect of a true competitive threat is now real and efficiency gains over the next decade could be larger.

Further reading

Ahlin, Å (2003) 'Does School Competition Matter? Effects of a Large-scale School Choice Reform on Student Performance', Department of Economics, Uppsala University Working Paper No. 2.

Björklund, A, M Clark, P-A Edin, P Fredriksson and A Krueger (2005) *The Market Comes to Education in Sweden: An Evaluation of Sweden's Surprising School Reforms*, Russell Sage Foundation.

Björklund, A, P-A Edin, P Fredriksson and A Krueger (2004) 'Education, Equality and Efficiency – An Analysis of Swedish School Reforms during the 1990s', IFAU Report No. 1.

Böhlmark, A and M Lindahl (2007) 'The Impact of School Choice on Pupil Achievement, Segregation and Costs: Swedish Evidence', IZA Discussion Paper No. 2786.

Böhlmark, A and M Lindahl (2008) 'Does School Privatization Improve Educational Achievement? Evidence from Sweden's Voucher Reform', IZA Discussion Paper No. 3691.

Bunar, N (2009) 'Can Multicultural Urban Schools in Sweden Survive the Freedom of Choice Policy?', Stockholm University Linnaeus Center for Integration Studies Working Paper No. 3.

Sandström, FM and F Bergström (2005) 'School Vouchers in Practice: Competition Will Not Hurt You', *Journal of Public Economics* 89: 351-80.

A silver lining to natural disasters:

what happened to student evacuees from Hurricane Katrina?

What would happen to educational achievement if failing schools were closed down and their students sent elsewhere? *Bruce Sacerdote* offers a glimpse of the possible impact in his research on the outcomes for student evacuees from Hurricane Katrina and for students in the schools where the evacuees ended up.

US Secretary of Education Arne Duncan made headlines in January with his public comment that Hurricane Katrina was 'the best thing that happened to the educational system in New Orleans'.

While Duncan's statement is controversial, there are many who believe that certain US school systems – such as Washington DC, Chicago and Los Angeles – are so broken that the only way to improve outcomes is to scrap them and either form new traditional public schools or allow charter schools to serve a large part of the market. Furthermore, a key component of the Bush-era 'No Child Left Behind' legislation is that failing schools that do not make annual progress will be shut down or reconstituted.

Hurricanes Katrina and Rita provide a glimpse into what might occur if entire failing school systems are shut down and the students sent elsewhere. My research studies the effects for the student evacuees themselves and the peer effects for 'native' students in the receiving schools.

Hurricane Katrina struck New Orleans in September 2005. The storm and resulting flooding killed nearly 2,000 people and caused at least $81 billion in damage. The insurance payments plus subsequent federal spending to mitigate the effects of the storm total in the hundreds of billions of dollars.

The experience of Katrina evacuees shows that leaving a very poor school system can significantly raise academic achievement

The vast majority of the damage and fatalities was caused by flooding that occurred when several levees in New Orleans were breached. Most schools in Orleans Parish (which is coterminous with the City of New Orleans) were flooded and closed for the year. Many schools in surrounding parishes – Jefferson, St Tammany, Plaquemines and St Bernard – were also closed for several months.

Roughly 66,000 students were evacuated from New Orleans. One third of these students left the state permanently. Another third relocated permanently to other districts in the state, including the East Baton Rouge Parish school district. The remaining third eventually returned to New Orleans after a period of six to 24 months.

My research suggests that Hurricane Katrina did indeed improve long-run outcomes for students from New Orleans, though perhaps in a different way than Secretary Duncan had in mind. From my analysis of data on test scores for all Louisiana students from Spring 2000 to Spring 2009, it seems that students who left New Orleans for good benefited the most.

Initially both Hurricanes Katrina and Rita had a large negative effect on maths and English scores for elementary, middle and high school students. This can be seen in Figures 1 and 2, which chart test scores for various groups of evacuees over time.

I standardise all test scores at the year and grade level. Hence Figures 1 and 2 show the performance of evacuees *relative* to the performance of all other students in Louisiana. The numbers in Figures 1 and 2 are coefficients from a set of regressions in which I calculate the coefficient on 'evacuee' controlling for race, gender and free lunch status.

Controlling for these demographics, the (eventual) evacuees from Orleans Parish are underperforming other Louisiana students by roughly .20 standard deviations in the pre-hurricane period. Following the hurricane, the performance of evacuees from Orleans Parish worsens by .12 standard deviations in Spring 2006. This negative effect of .12 standard deviations seems to be consistent across different groups of evacuees and across the two different hurricanes.

But one year later, the Orleans evacuees have made up the ground they lost due to the hurricane and associated

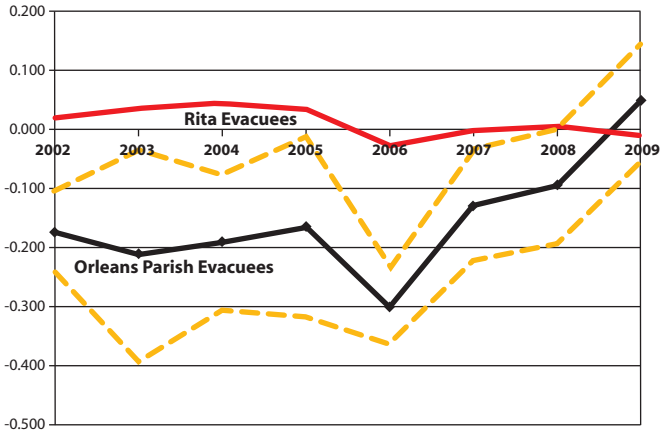

Figure 1: Repeated cross-sectional 'effects' on maths scores, New Orleans versus non-New Orleans evacuees

I regress maths scores (all grades) on dummies for eventual Katrina Evacuee Status. The latter is split by evacuees who are in Orleans Parish in 2004 or 2005 versus all others. The 2006 and 2007 scores are post-hurricane. Students are tested in March of each year.

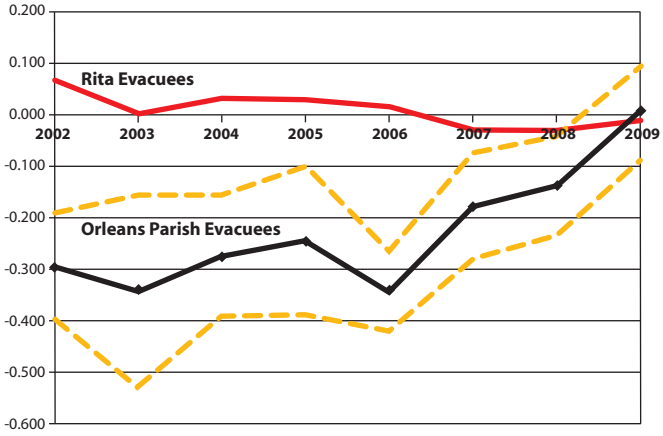

Figure 2: Repeated cross-sectional 'effects' on ELA scores, New Orleans versus Rita evacuees

I regress English Language Arts scores (all grades) on dummies for eventual Orleans and Rita Evacuee Status, race dummies, male, free lunch status. The latter is split by evacuees who are in Orleans Parish in 2004 or 2005 versus all others. The 2006 and 2007 scores are post-hurricane. Students are tested in March of each year.

disruption. In the subsequent two years (2008 and 2009), the Orleans evacuees have made substantial progress. By 2009, the evacuees are performing .20 to .25 standard deviations above where they started.

There are several possible explanations for this surprising finding. One sensible explanation is that the Orleans schools were so ineffective that the switch in schools had long-term positive effects on test scores. As one check of this hypothesis, I examine test score growth for evacuees conditional on where the evacuees settled.

Students with the lowest prior attainment and those who left the metro area entirely saw the greatest improvements

Evacuees who left the New Orleans metro area have much higher test score growth from 2005 to 2009 than evacuees who return. Those evacuees who return to the city of New Orleans itself and enter the Recovery School District do not experience any noticeable growth in their position in the test score distribution. In other words, test score growth for returning Orleans students is comparable to that experienced by all other Louisiana students.

A second question is whether the improvement in outcomes experienced by Orleans evacuees is concentrated in one part of the achievement distribution. Figure 3 addresses this question, showing the amount of test score growth from 2000 to 2009 stratified by initial quintile of achievement (test scores).

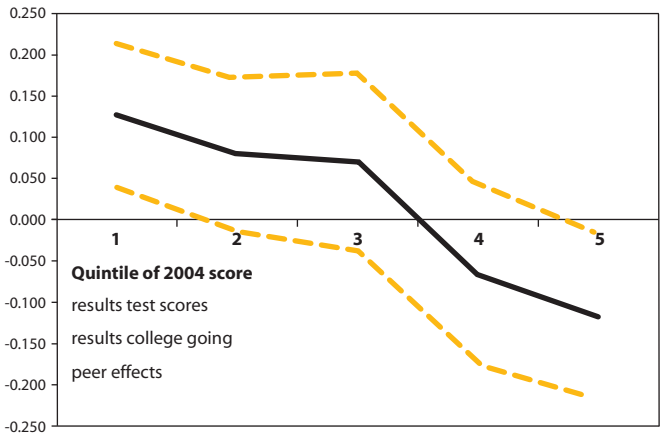

Figure 3: Effect by quintile of initial maths scores, Orleans 4th graders of 2004 in 2009

This runs the baseline specification by quintile (statewide) of initial maths score in 2004. I graph the coefficient of Orleans Evacuees 2009.

I find that evacuees who were in the bottom two quintiles of the Louisiana test score distribution benefit the most from the disruption induced by the hurricane. This suggests that the students who needed the most help received that help once they left the poorly performing Louisiana school system. (Again these effects are all relative to other students in Louisiana in the same year and grade. Thus a coefficient of 0 suggests growth that is in line with all other students.)

While the figures are a very simple way of examining the data, my research also makes use of more sophisticated specifications in which I am careful to limit the sample to particular cohorts of students who are observed throughout the sample period (for example, the fourth graders from 2004 and 2005.)

Test scores are only one outcome measure of interest. I also have data on eventual college attendance for every single Louisiana tenth grader from 2000 to 2005. When I analyse college attendance patterns for the evacuees, a different picture emerges than the results for test scores.

The hurricane had virtually no impact on the likelihood that an evacuee from Orleans enrols in a four-year college. One possible explanation is the fact that the test score effects take place in the lower part of the test score distribution and these students are the least likely to attend college, even with their improved scores.

When I consider effects on eventual enrolment in two-year (community) colleges, the hurricane had negative effects on enrolment. Students attending an Orleans Parish high school at the time the hurricanes strike are 3.5 percentage points less likely to enrol in community college compared with students who attended those same Orleans high schools but graduated one to three years before the hurricanes struck.

Schools where the evacuees went were able to absorb the inflow without causing much harm to their own students

The explanation for this finding is that there are very close 'feeder' relationships between high schools in a particular area and the local community colleges. Before the storm, two thirds of students from Orleans who attended a community college attended Delgado Community College. The evacuees' exit from New Orleans appears to have disrupted the natural flow of students from their high school to the local college.

Finally, in research conducted jointly with Scott Imberman and Adriana Kugler, I investigate whether students in the receiving schools see large effects from the influx of evacuees. On average, it appears that the receiving schools (in Houston, Texas, and East Baton Rouge and Shreveport, Louisiana) were able to absorb the inflow of evacuees without significant negative effects on their native students.

High- and moderately high-achieving 'native' students were helped by the arrival of high-achieving evacuees

But within this average effect of zero, there are large and interesting nonlinearities. In particular, low-achieving native students are harmed by the arrival of low-achieving evacuees. High- and moderately high-achieving 'natives' are helped by the arrival of high-achieving evacuees.

Overall this work suggests that we may able to learn from studying students who are forced to relocate due to a natural disaster. The experience of Katrina evacuees from New Orleans shows that leaving a very poor school system, such as the Orleans Parish School district, and entering a more typical US school district significantly raised academic achievement for those students.

The students who had the lowest baseline scores and who left the metro area entirely saw the greatest gains. And these gains were generated without causing any harm (on average) to students in the receiving schools. But there does not appear to be a positive effect on college attendance and, if anything, transitions from high school to community college seem to be hindered.

This article summarises 'When the Saints Go Marching Out: Can Hurricanes Katrina and Rita Teach us about Effects from Closing Unsuccessful Schools?' by Bruce Sacerdote of Dartmouth College.

School choice: what do parents want?

Do different types of parents have different preferences about the type of school they want their children to attend? *Deborah Wilson* and colleagues explore what parents say they want – and which schools they actually choose.

The merits of different schools, and strategies for getting children a place at the most desirable ones, are famously the focus of countless dinner parties. But what are parents actually looking for in a school? This question matters not only for understanding the process by which children are assigned to schools, but also for evaluating the impact of school choice on educational outcomes.

The theory is that if families care most about academic standards, then competition between schools will focus on this and encourage schools to raise standards overall. But if parents value other characteristics in a school, such as its social or ethnic composition, the result of school choice may be increased social stratification. Parents' preferences for different types of schools with different characteristics fundamentally affect the outcome of any choice-based assignment mechanism.

All parents have a strong preference for schools with high academic attainment

A joint project by CMPO and the Institute of Education has been investigating what parents want from schools. Specifically, we examine whether different types of parents have different preferences about the type of school they want their children to attend.

A key question is whether more socially advantaged parents place greater emphasis on academic standards. If so, school choice would tend to increase social segregation across schools, with advantaged pupils concentrated in schools with the highest academic standards.

Our focus is on primary school choice, and we use the Millennium Cohort Study (MCS), a nationally representative dataset that is tracking nearly 19,000 children born in Britain around the turn of the century. The third wave of this survey, carried out in 2006 after the children had started school, includes detailed questions on the school choice process, including the reasons parents state were important in their choice, whether they achieved their first choice and whether they acted strategically in trying to do so.

The survey data also include details of up to three schools nominated by parents on their local authority school application form, as well as detailed information about the background of the family and the child.

We are therefore able to examine both what parents say they want from a school (their 'stated preference') and what they actually revealed they wanted from a school as indicated by the characteristics of the school they put first on their local authority form (their 'revealed preference'). In practice, when filling in the form, parents are choosing between a range of 'local' schools.

By combining the MCS with two national administrative school datasets (the Pupil Level Annual Schools Census/National Pupil Database and Edubase), we are able to work out which schools are available for each child and compare the parents' first choice school with others in the local area. This gives us an idea of what types of schools parents prefer from the options available to them.

Obviously, some schools have small and restrictive catchment areas, and therefore we cannot simply look at pupils' nearest schools. Rather, we use a method that enables us to determine which schools were likely to be really available to them.

Our analysis of what parents stated were their preferences for schools is in line with previous research. We find that parents who are more educated and of higher socio-economic status are more likely to say that they value academic standards above

other factors. Parents who are less educated and of lower socio-economic status are more likely to cite the school's proximity to their home as the key factor.

Stated preferences, however, may not tell the whole story. Parents may only give socially desirable responses or they may forget some of their initial considerations. It is also hard to understand the trade-offs that families face between different characteristics of schools just by using data on what parents say are their preferences.

Differences in location rather than parents' preferences drive differential access to higher-performing schools across socio-economic groups

That is why we also look at the schools that parents actually choose, relative to the other schools available to them. We think such revealed preferences are more informative about what parents are really looking for from schools.

In doing this, we need to take account of the fact that within a given area, a family may not be able to gain admission to a particular school, even if it is quite close to their home. It may be that while a pupil lives within a given distance of five or six schools, their address will influence which of these schools they can actually get into. For example, poorer parents are likely to live much closer to schools with higher proportions of children eligible for free school meals.

When we take account of this restricted choice, we confirm that all families do indeed choose schools on the basis of their academic performance, as measured by the percentage of pupils exceeding the expected level at key stage 2. Parents also value particular peer groups, preferring schools with low proportions of poor children.

All parents value proximity and they trade this off against their preferences for academic performance and peer group. We find that on average a family will be willing to travel another 400 metres or so to reach a school with higher attainment of 11 percentage points.

We also find that richer and poorer families actually have similar preferences across school characteristics. So allowing for the fact that poorer families tend only to be able to access relatively lower-performing schools with higher proportions of poor children, we find that within this limited group of schools, richer and poorer parents make similar choices. It is differences in location rather than preferences that drive differential pupil assignment across schools.

What of the annual statistics published on the numbers of parents who do not achieve their first choice school? Actually, in our data, 94% of families gain a place at their first choice, and this is true for advantaged and disadvantaged families.

We do find, however, that this high proportion is partly explained because less advantaged parents make more 'realistic', less ambitious choices of schools. In other words, although the high proportion of parents getting their first choice of school makes it seem as though school choice is working effectively, in fact many parents, particularly disadvantaged ones, recognise that they are unlikely to get into some schools and therefore make less ambitious choices in the first place.

Looking forward, it may be that as parents get more used to the new school admissions rules or as other 'tie-breaker' schemes (such as ballots) are introduced, more parents will not achieve their first choice. This in turn may put more pressure on the system to expand the supply of places in high-performing schools.

So what are the implications of our findings for educational inequality and for policies on school choice? On educational inequality, we show big differences in the choices available to different families. We show less important differences in their preferences.

The big driver of differential access to higher-performing schools is the quality of schools near to where families live. This is because the main factor in determining whether a child can get into an oversubscribed school is their geographical proximity to that school. This research therefore relates immediately to practical issues about the operation of the current system in England, particularly school admission procedures.

The huge differences in the types of schools that rich and poor families can access illustrate the difficulty of reducing educational inequality

The broader implications of our results for choice in education are mixed. Almost universally in our data, parents have a strong preference for schools with high academic attainment. This supports the idea that competition to meet those preferences should help to raise standards.

The huge differences in the types of schools that rich and poor families can access are an illustration of the problems involved in reducing inequalities of access. And the importance of preferences for peer groups means that there is also pressure for socio-economic 'sorting'. The challenge for education policy is to address the former while minimising the potential for the latter.

This article summarises 'What Parents Want: School Preferences and School Choice' by Simon Burgess, Ellen Greaves, Anna Vignoles and Deborah Wilson, CMPO Working Paper No. 09/222 (http://www.bris.ac.uk/cmpo/publications/papers/2009/wp222.pdf).

Politics and tax levels:
evidence from American states

What is more important in determining the level of taxes in individual American states – political institutions or political ideology? *Leandro de Magalhaes* and *Lucas Ferrero* look at nearly half a century of data to explore this question.

Each of the 50 states in America works in a similar way to the US federal government: it has both an executive and a legislature. Each state can also create state-specific taxes on top of federal taxes. To understand what determines the level of taxes in American states, we analyse data on state budgets, political institutions and election outcomes covering the period from 1960 to 2006.

Our measure of a state's tax level is the sum of state-specific income, corporate and sales taxes divided by the state's GDP in that year. And the first question we ask is whether Democratic political control of a state has a causal effect on the tax level. As Figure 1 shows, there is a positive relationship between the tax level and the percentage of seats the Democrats hold in the state legislature. This is in line with our expectation that Democrats will typically prefer a bigger government and higher taxes.

But what does this relationship mean? We cannot conclude from a positive correlation that there is a causal link between Democratic control and a higher tax level. In other words, we cannot conclude that taxes would go up if we were to keep everything else constant and increase the number of Democrats in the state legislature.

Tax levels in American states are not driven by whether Democrats or Republicans control the legislature

It may be the case that voters' preference for a certain tax level varies for reasons that have nothing to do with ideology. Changes in the public's preferred level of taxation could be driven by other variables, such as unemployment levels, a crisis in a specific economic sector, immigration to the state or other unknown variables. Representatives in the state legislature may just be responding to these preferences.

If this is indeed the case, when voters start to demand higher taxes, the representatives will deliver it regardless of whether they are Republicans or Democrats. What Figure 1 captures could just be that when the voters prefer higher taxes, they also tend to vote Democrat. What drives the tax level, however, is not the degree of Democratic political control but the variables that determine voters' preferences. Representatives do just what voters want, be they Democrats or Republicans.

Tax levels are influenced by whether executives and legislatures are politically aligned or politically divided

To check whether Democrats have a causal effect on the tax level, we restrict our attention to a small variation in the percentage of Democrats in the legislature that leads to a major shift in political control. We look at what happens when the percentage of Democrats in the state legislature crosses the 50% threshold. If the Democrats have more than 50% of the seats, they can impose their preferred budget and tax level on the Republican minority. If they have less than 50%, the Republicans can impose their preferred budget and tax level on the Democrats.

We compare the tax level between elections that have delivered slim majorities, as we can think of election results around the 50% threshold as random. These elections were so close that the outcome (which party won the majority) was decided by random events such as rainfall and its effect on turnout.

If we accept this assumption, then all the variables that may

Figure 1: State tax level and *Democratic control*

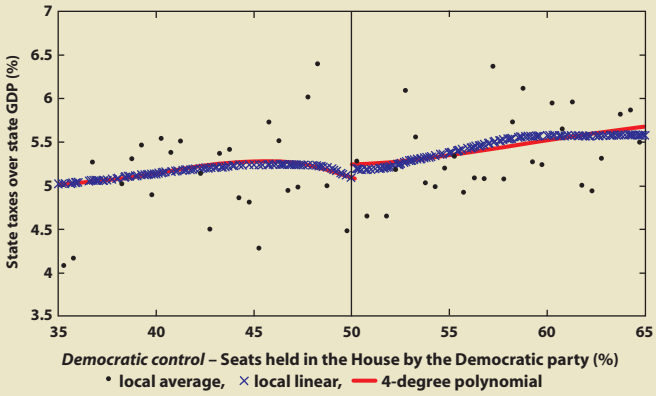

Figure 2: State tax level and *Governor's strength* in states with line-item veto

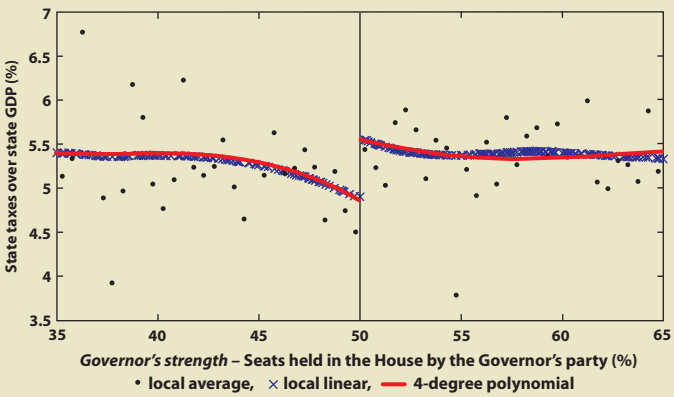

affect voters' preferences, such as unemployment, economic conditions, immigration and others, must, on average be the same on either side of the 50% threshold. And if they are the same on either side, they could not be causing a change in the tax level at the 50% threshold.

These random election results may exchange a few Republicans seats with a few Democratic seats, turning a slim Republican majority into a slim Democratic majority. If this happens, and the tax level increases, it must be because the Democrats have gained the majority, as this is the only variable that has changed as we moved across the 50% threshold.

In Figure 1, we observe no jump in the tax level at the 50% threshold. This is so even though we estimate the average tax level immediately to the left and right of the 50% cut-off only using data to the left and right respectively. This result suggests that adding Democrats to the legislature does not have a causal effect on the tax level. The party identity of the majority, it seems, has no causal effect on the tax level.

We then explore an alternative hypothesis. Most states give the governor more power over the state budget than the US president has over the federal budget. In most states, the governor can veto or trim particular items or lines in the budget. In a few other states and in the federal government, the executive can only block veto the budget – it cannot selectively cut the budget.

If the objectives of the governor and of the majority party in the legislature are aligned – that is, they are both from the same party – there should be little to cut. If, however, there is a divided government – that is, the majority in the legislature is from one party and the governor is from another party – we would expect the governor to take advantage of the line-item veto power.

To test this hypothesis empirically, we look at the percentage of seats in the state legislature belonging to the governor's party, whether Democratic or Republican. We call this variable *governor's strength*. Again, as we move across the 50% threshold, there is a major shift in political control. To the left of the threshold, the government is divided, whereas to the right of the threshold, the government is aligned.

Figure 2 shows that as we move across the 50% threshold, the tax level jumps. This indicates a causal relationship between having an aligned government and a higher tax level. We estimate that the tax level increases by 13%.

Figure 3: State tax level and *Governor's strength* in states with block veto

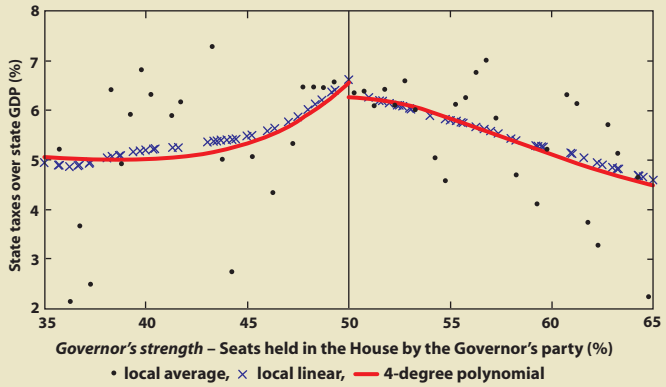

Interestingly, we only observe this jump in states in which the governor has the line-item veto. In states in which the governor has the block veto, there is no jump in the tax level, as Figure 3 shows.

Our theoretical explanation for this is that the budget can be thought of as a 'sequential bargaining game' between the majority party in the legislature and the governor. The legislature makes an offer with different spending items and an overall tax level. In states with the block veto, this is a 'take-it-or-leave-it' offer.

It is too costly for the governor to block veto this offer. In most states, blocking the whole budget implies a government shutdown. Governors are therefore reluctant to use their block veto power. In these states, the majority party in the legislature is able to target expenditures as it wishes. This gives the majority an incentive to have high taxes, as they benefit fully from an extra dollar of revenue.

The design of a state governor's budget veto – block veto or line-item veto – has an impact on the tax level

In states with the line-item veto, the budget is not a 'take-it-or-leave-it' offer. Once the budget has been approved by the legislature, the governor can selectively cut particular items or trim down values. In these states, the majority does not have the freedom to target spending to their benefit. A governor from another party would easily veto any item. Since the majority cannot choose how to spend the tax revenues, they do not have any incentives to propose a high tax level.

Our three results seem to indicate that party identity – that is, whether the majority is Democratic or Republican – is not the key determinant of the tax level. Rather, it seems that it is institutional features – the type of budget veto power and whether the government is aligned or divided – that are the determining factors.

This article summarises 'Budgetary Separation of Powers in the American States and the Tax Level: A Regression Discontinuity Design' by Leandro de Magalhaes and Lucas Ferrero, CMPO Working Paper No. 09/225 (http://www.bristol.ac.uk/cmpo/publications/papers/2009/wp225.pdf).

Measuring diversity in England's schools: a new web resource

A new online resource from CMPO sheds light on ethnic segregation in England's schools. *Simon Burgess* introduces the website.

One of the biggest questions of our times is how well individuals from different ethnic groups get along together. Schools are an important place where this interaction takes place. It is a common saying that people's attitudes are strongly influenced by their school days. So the peer groups that children play with, talk to and work with are important factors moulding their perspectives on society.

The extent of ethnic diversity in schools is an important issue of public debate. Our new website – www.measuringdiversity.org.uk – tries to enlighten this debate by providing up-to-date and comprehensive information on ethnic segregation in schools in England. Users will also be able to follow trends over the past seven years, up to 2008, the most recently available data.

For example, the website shows that:

- One in eight (13%) secondary school pupils in Manchester are of Pakistani ethnicity. Fewer than one in a hundred (1%) of these pupils attend schools that have mostly white pupils, while over half (55%) attend schools that are mostly non-white. At the same time, segregation of pupils of Pakistani ethnicity in Manchester has been falling over the last seven years, 2002 to 2008.

- Pakistani and Bangladeshi pupils in Oldham account for around 30% of the primary school population. Around 80% of both Pakistani and Bangladeshi primary school pupils are in schools that are mostly non-white and this proportion has remained roughly constant over time. The 'dissimilarity index of segregation' for Oldham is very high at around 0.8 for all ethnic groups: this means that 80% of pupils would have to change schools to achieve the same distribution of pupils in schools to that of the local authority.

- The most segregated ethnic group in Camden are Bangladeshi pupils. Three quarters (76%) of Bangladeshi primary school pupils attend schools that have mostly non-white pupils, compared with just under half (49%) of Black African pupils and one in six (17%) white pupils.

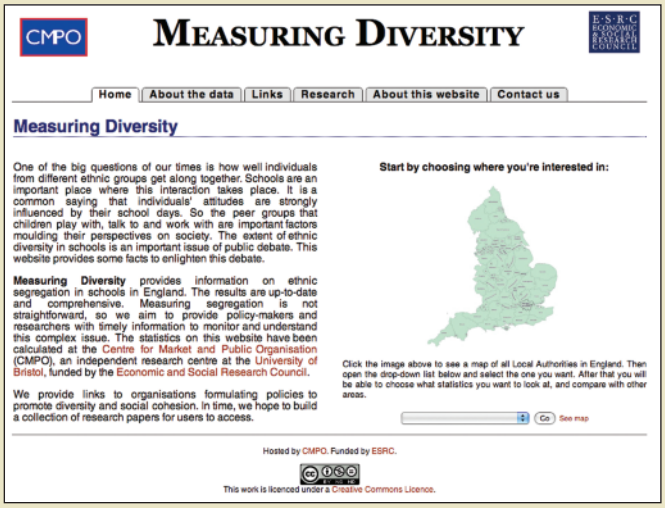

- The percentage of white pupils at primary schools in Wolverhampton has gradually decreased from 69% in 2002 to 57% in 2008, although the level of segregation has remained roughly constant.

Measuring segregation is not straightforward, so the aim is to provide policy-makers and researchers with timely information to monitor and understand this complex issue. Users of the website start by picking a local authority from a map or a list, and choosing what statistics to view. Results are presented for the most numerous individual ethnic groups across England. The available statistics include different measures of the ethnic composition of schools and different measures of segregation.

All of the statistics are fully explained in plain English. Users of the website can choose between viewing details for an individual year or trends over time. Users can also pick another area to compare with their chosen local authority, and can choose to make the comparison for all or just selected ethnic groups. Tables and graphs are downloadable.

The website – www.measuringdiversity.org.uk – was launched in January 2010. It was created by Simon Burgess, Ellen Greaves and Simon Speight.

Research evidence on early life health

A growing body of economic research is exploring the influence of pre-natal and early childhood conditions on later life outcomes and future generations. A recent CMPO workshop, organised by *Sonia Bhalotra*, discussed some of the latest findings.

There is growing evidence of the long-term benefits of 'investing' in children's health, and recognition that investments made during critical periods of their development draw larger returns – and, conversely, that failure to invest can lead to irreversible damage. The origins of life in the womb constitute one such period, when growth is rapid and neurological and physiological development is particularly sensitive to the environment.

Nutritional deprivation, for example, appears to lead to permanent changes in tissue structure and function that help the foetus survive, adaptively, through pregnancy. At the same time poor nutrition *in utero* exacts a penalty in terms of later life outcomes such as, diabetes and cardiovascular disease, reduced adult height, obesity among women and lower cognitive outcomes and earnings.

Some studies have directly analysed sharp disruptions in mothers' nutrition as a result of famine. The Great Chinese Famine of 1958-61 and the Dutch famine of 1944, known as the 'hunger winter', are examples. Others have used recessions in pre-industrial settings or recessions and infant mortality rates in developing countries as proxies for nutritional 'shocks'.

Further studies have documented long-run consequences of exposure of pregnant mothers to infection and radiation. Studies of the impact of smoking on birthweight naturally blend into this research literature too, smoking being a form of pollution and birthweight an often-used proxy for longer-term effects.

The studies presented at the recent CMPO workshop offered important new contributions to this rapidly advancing field. Most are summarised by the authors in the following pages; below is an overview.

Health shocks during pregnancy

A paper by Elaine Kelly explores the consequences of *in utero* exposure to the Asian influenza pandemic which struck Britain in 1957-58. The findings indicate significant detrimental effects of the epidemic on intrauterine growth and, conditional on birthweight, on test scores at 7 and 11. These effects are either restricted to or more pronounced among women who are relatively short or who smoke during pregnancy.

Kelly's findings reinforce earlier research by Douglas Almond, which documents the lasting impacts of the Spanish flu pandemic of 1918 in the United States.

Almond is co-author of the next paper (with Bhaskar Mazumder), which investigates the effects of the nutritional shock of maternal fasting during pregnancy. Thousands of Muslim women invite nutritional disruption by fasting through the month of Ramadan each year. This study finds that pre-natal exposure to Ramadan results in lower birthweight, reduced gestation length, a reduced probability that the birth is male, increased likelihood of learning disabilities in adulthood and adverse effects on schooling, earnings and measures of wealth.

Environmental conditions in childhood and intergenerational effects

Nutritional deprivation in childhood is more closely tied to aggregate income in poorer countries. A study by Gerard van den Berg and colleagues set in 19th century Denmark shows that individuals born in recessionary conditions are more likely to suffer cardiovascular disease in later life but they appear to be no more likely to suffer from cancer. The effects of genetic factors and family background are stronger when early life conditions are poor.

Sonia Bhalotra and Samantha Rawlings suggest a similar finding using data for 38 developing countries. The influence of maternal health on child health is stronger when early life conditions are poorer. Maternal health is indicated by height, which incorporates genetic and family background influences. Since adult height is an indicator of the scars of environmental deprivation in childhood, this suggests an intergenerational impact of early childhood conditions.

Bhalotra's work in progress establishes this directly, showing effects of nutrition and infectious disease in the mother's birth year on the survival chances of her children. These effects hold conditional on the mother's height. This indicates that while height captures some of the scars, mothers are scarred in ways that are not captured in height and which have an impact on the next generation.

As long ago as the 1920s, scientists cautioned against adding a known poison to petrol, highlighting the effects of lead on birthweight, sperm abnormalities, neurological development and the risk of stroke and heart disease. Four decades later, a series of regulations were introduced across the United States and Europe in the early 1970s with a view to phasing out lead in petrol.

Peter Nilsson analyses the impact of the phase-out in Sweden, which ran from the late 1970s through to the mid-1990s. He finds that reduced exposure early in life improves cognitive performance and labour market outcomes among young adults, and more so among families of low socio-economic status.

This issue remains topical. Lead pollution is now concentrated in developing countries although rapid phasing out is currently in progress. Initially, US corporations exported their lead additives overseas, using the profits to finance diversification at home.

More recently, international donors and corporations are helping developing countries achieve a phasing out. Britain has not completed its phasing out of lead in petrol, and there remains lead exposure from sources other than petrol such as lead-based paint, pipes or toys in many developed countries.

The long-term benefits of investments in foetal and early childhood health

Together, these studies present compelling evidence to suggest that the benefits of investments in foetal and early childhood health may extend into adulthood and, through mothers, into the next generation.

To what extent can pregnant mothers function as an effective buffer for the foetus against poor environmental conditions? The buffering capacity of mothers may explain why the effects in some studies are small or insignificant. It probably explains the heterogeneous effects documented by Kelly, whereby children of short mothers or mothers who smoke during pregnancy are most affected by the mother contracting influenza.

It may similarly explain Bhalotra and Rawlings' finding of a stronger impact of mother's height on neonatal mortality among mothers of relatively short stature and, for a given maternal stature, the greater sensitivity of neonatal mortality to shocks to income or public health among relatively short women.

The significance of maternal height indicates that it is not only the contemporaneous health of the mother or health during pregnancy that counts but rather the stock of her health. In view of evidence that adult height is sensitive to the early life health environment, this brings into focus intergenerational impacts flowing from the early life of the mother.

The scourge of Asian flu: effects of *in utero* exposure

In the late 1950s, Britain was struck by the Asian influenza pandemic. *Elaine Kelly* examines the consequences for children who were *in utero* at that time. Her findings are revealing about the channels by which health conditions experienced in the womb influence later life outcomes.

The foundations for lifelong health and acquisition of education and skills are laid in the womb. Economists have typically tried to understand the influence of 'the intrauterine environment' on later life outcomes by analysing circumstances in which there have been sharp and random shocks to foetal health conditions.

For example, work on the devastating Spanish flu of 1918-19 has established that exposure to pandemic influenza is one such potential shock. Almond (2006) finds that people exposed during the first or second trimester of their mothers' pregnancy were less likely to complete high school, had consistently lower earnings, received higher welfare payments and had higher rates of incarceration than would be expected.

My research examines the impact of intrauterine exposure to the Asian influenza pandemic of 1957-58 on outcomes at birth and in childhood (Kelly, 2009). Asian flu began in China and although it was far milder than Spanish flu, between one and two million perished worldwide.

When the pandemic struck Britain in the autumn of 1957, the children in the National Child Development Study (NCDS) were at between 15 and 25 weeks in gestation. The NCDS follows a panel of over 17,000 individuals born in Britain in one week in March 1958. My study uses data from the first three waves of data gathering on these individuals, conducted in 1958, 1965 and 1969.

Outcomes at birth are measured by birthweight and gestation. Birthweight is used for two purposes: first, as an indicator of whether the effects of the epidemic can be detected at birth; and second, to assess whether outcomes measured at birth capture the effects of foetal health shocks on subsequent development. At ages 7 and 11, the outcomes used are cognitive test scores and height in metres.

The epidemic is measured using a surrogate influenza infection rate in the cohort members' local authority of birth, namely, official pneumonia notifications per 100,000 population (see maps). The impact of Asian flu on outcomes is identified using variation in the intensity of the epidemic measure across the 172 local authorities. All estimations include controls for children's background and for the health and socio-economic characteristics of the local authorities.

The study adds to the research on Spanish flu by providing results for outcomes at birth and in early childhood. But the more substantive contributions stem from the panel structure of the NCDS.

Exposure to Asian flu had a significant and negative effect on birth outcomes for some children

Two issues are addressed in detail: first, the extent to which the effects of the epidemic on childhood outcomes are captured by birthweight; and second, the role of maternal health in moderating the effects of influenza. The principal aim is to gain some information on the mechanisms that might link foetal health shocks to subsequent outcomes, and the findings are as follows.

Birth outcomes

Asian influenza had a significant and negative effect on birth outcomes, but only for the offspring of mothers with certain health characteristics. There was no effect of the epidemic on average birthweight, gestation or intrauterine growth (birthweight corrected for gestation).

But analysis of interactions between the epidemic and two indicators of maternal health – pre-pregnancy smoking and adult height – shows that negative effects on birth outcomes were confined to particular groups of women: those that smoked before pregnancy; and those who were 61 inches (155 cm) or shorter.

Smoking weakens the respiratory system, increasing the risk of infection and the severity of symptoms. Height is an indicator of

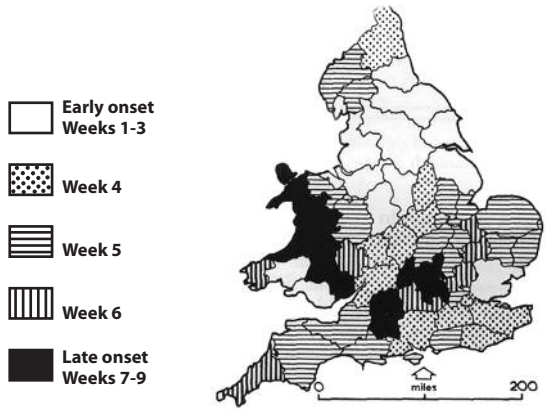

Figure A: Spatial timing of the epidemic

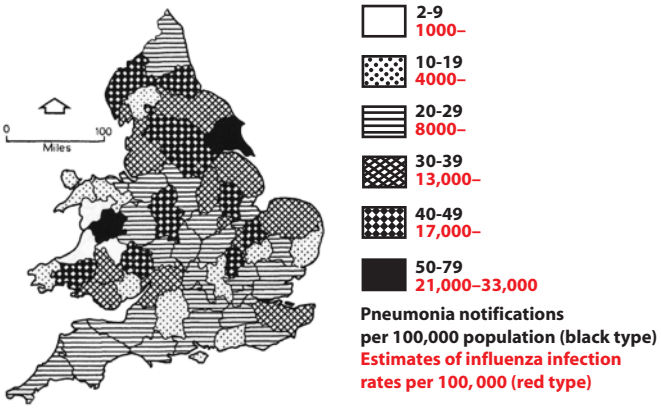

Figure B: Total incidence of the epidemic

the long-run nutritional status of the mother, and therefore the store of nutrients from which the baby can draw.

Outcomes at ages 7 and 11

The epidemic had significant effects on average test scores at age 7 and 11: a one standard deviation increase in epidemic intensity reduces test scores by 0.06 standard deviations at 7 and 0.05 at 11. By way of comparison, Rockoff (2004) finds that a one standard deviation increase in teacher quality increases reading and maths test scores by 0.1 standard deviations.

These effects are general, and do not vary with cohort member characteristics, including maternal health, socio-economic background and measures of parental investment.

The epidemic also had negative effects on children's height at 7 and 11 but only for the children of smokers.

Does birthweight mediate the impact of the epidemic on test scores?

The effects of the epidemic on childhood height and test scores do not operate through, and are not captured by, birthweight. The association between birthweight and later physical and cognitive outcomes is well documented by medical and social science research, but it remains poorly understood. My results suggest that birthweight should not be viewed as a catchall measure of influences on development prior to birth.

Transmission mechanisms

The effects of Asian flu appear to have operated through two distinct channels. Physical development growth is impeded where mothers are unable to compensate for interruptions in nutrition (indicated by maternal height) or when symptoms are plausibly more severe (through maternal smoking). The effects on cognitive development are more general, and may be related to the inflammation that typically accompanies a flu infection.

The implications of these results are threefold:

First, the effects of foetal health shocks are not necessarily observable at birth. In particular, there may be no effect on an individual's birthweight or gestation. It may therefore take years or decades before the full ramifications of a particular health shock become clear.

Second, the effects of a foetal health shock might vary considerably across those who are exposed.

Third, maternal smoking plays the most consistent role in moderating the effect of the epidemic on physical growth. The negative interaction between smoking and influenza emerges at birth, and continues into adulthood. Whether these results are attributable to smoking during pregnancy or to the damage caused by previous smoking behaviour, it is not possible to determine using these data. Women are already advised not to smoke during pregnancy. My results perhaps provide one more reason to give up.

Elaine Kelly is at University College London and the Institute for Fiscal Studies (IFS).

Further reading

Almond, D (2006) 'Is the 1918 Influenza Pandemic Over? Long-term Effects of in utero Influenza Exposure in the Post-1940 US Population', *Journal of Political Economy* 114(4).

Kelly, E (2009) 'The Scourge of Asian Flu: in utero Exposure to Pandemic Influenza and the Development of a Cohort of British Children', IFS Working Paper No. 09/17.

Rockoff, JH (2004) 'The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data', *American Economic Review* 94(2): 247-52.

The intergenerational transmission of health in developing countries

A harsh indicator of poverty and inequality in poor countries is the prevalence of under-nutrition, which combines with infectious disease to produce widespread morbidity and mortality. Scarce resources are often allocated preferentially to sons even in countries without a marked preference for sons, possibly because the pecuniary return to investment in sons is perceived to be higher, whether because of their greater labour market attachment or the prerogatives that property rights confer on them.

Our research suggests that this behaviour may be myopic since investments in the nutrition and health of women have an intergenerational payoff, improving the survival and health of future generations.

Previous studies have established the importance of good health in raising educational attainment and productivity, and of poor health in generating poverty traps. Together with evidence of the impact of education and income on health in poor households and of 'assortative matching' in marriage, this implies that intergenerational persistence in health may explain some of the intergenerational persistence in education, earnings and economic inequality.

There is a strong and widespread tendency for children born to relatively unhealthy mothers to start life in poor health

Our analysis uses comparable micro data on 2.24 million children born to about 0.6 million mothers in 38 developing countries in the 31 year period, 1970-2000. The research makes two main contributions. It presents the first estimates of the intergenerational transmission of health across countries and over time; and it investigates the extent to which environmental conditions around the birth of a child mediate the intergenerational transmission of health.

A mother's health is indicated in our research by her height, body mass index (BMI) and anaemia status. Child health is indicated by mortality risk and 'anthropometric failure', namely, low birthweight and stunted growth.

We find that short stature and anaemia of the mother each raise the likelihood that her children are low birthweight, suffer early life mortality and exhibit stunted growth. Low maternal BMI is a risk factor in low birthweight and stunting, while high BMI is a risk factor for childhood mortality. A one standard deviation decrease in mother's height or BMI raises the risk of poor child health by between 5% and 10%, depending on the outcome. The intergenerational transmission of health is a widespread phenomenon, statistically significant in between 21 and 29 of the 38 countries, depending on the measure.

The intergenerational transmission of health appears to be weakened by improvements in income, maternal education and public health provision

There is stronger intergenerational persistence at the low end of the mother's health distribution. We may therefore expect average persistence to decline with improvements in maternal health. Alternatively, positive trends in public health programmes that effectively target children at most risk will tend to weaken persistence.

Investigating this for the 31 cohorts in 1970-2000, we find an erosion of about 20-30% per decade in the link between child mortality and mother's height. This looks fairly impressive, especially in view of the record of limited income and education mobility. But disaggregation by continent reveals that Latin America alone exhibits a consistent improvement in health mobility through the period, with rates of decline twice as large as the average rates. Asia shows no significant trend and Africa shows a worsening trend, especially for neonatal mortality.

These results indicate that children born to relatively unhealthy mothers in relatively poor regions in Africa and Asia start life in

To what extent are children's health outcomes determined by their mothers' health? *Sonia Bhalotra* and *Samantha Rawlings* analyse three decades' worth of data on over two million children across 38 developing countries to explore how health is transmitted across generations – and how public policy can respond.

poor health. The penalty they incur has shown no tendency to fade in the last 30 years, a period in which children born to better off mothers in these and other regions have prospered. The natural question for public policy concerns the extent to which differences in health at birth between children can be narrowed, so that children from disadvantaged families start life with more equal opportunities.

Improvements in the socio-economic environment at birth generate greater payoffs in less healthy families

Exploiting within-continent heterogeneity in trends in income and health, we test the hypothesis that the decline in the intergenerational transmission of health between cohorts separated by two decades has been faster in countries that recorded positive growth than in countries that experienced stagnation or negative growth. This 'experiment' is possible with these data because of the sharp diversity of growth experiences of developing countries in the last three decades. We find that economic growth is associated with a weakening of the tie between mother and child health.

But it is unclear whether this benefit is attributable to income since long-range growth is potentially confounded with other changes, such as medical technological progress. We therefore investigate the sensitivity of the intergenerational transmission of health to annual changes in aggregate income, mother's education and public health (indicated by immunisation rates).

We find that improvements in mothers' education, immunisation rates and income in the country and cohort of the index child weaken the intergenerational transmission of health. The gains from improvements in immunisation rates are evenly distributed, but the gains from improvements in income and maternal education are greatest for children who are initially most disadvantaged by being born of relatively unhealthy mothers. These results also suggest that children are more likely to bear the scars of poor maternal health if they are conceived or born in adverse socio-economic conditions.

Our finding that maternal stature has a substantial negative influence on a range of measures of child health contributes to evidence that adult height is an indicator of health. The finding that adult height is especially sensitive to the early childhood environment, together with our finding that children of shorter mothers are more sensitive to changes in the socio-economic environment, suggests that the intergenerational transmission of health involves not only genomic but also non-genomic mechanisms.

Investments in the nutrition and health of women have an intergenerational payoff, improving the survival and health of future generations

This research contributes unique evidence on an important and under-studied aspect of persistent inequality in developing countries, where underdeveloped markets and states result in children often being unable to escape from the family circumstances into which they are born. It paints the first broad-brush picture of the persistence of health across generations, while also presenting continent- and country-specific estimates and evidence on how the transmission of health may be weakened by improvements in income, maternal education and public health provision.

This article summarises two studies by Sonia Bhalotra and Samantha Rawlings: 'Gradients of the Intergenerational Transmission of Health in Developing Countries' and 'Intergenerational Persistence in Health in Developing Countries: Trends and Country Differences' (http://www.efm.bris.ac.uk/ecsrb/bhalotra.htm).

Air quality in early childhood: the impact on later life outcomes

From the end of the 1960s, government regulations on air pollution have become increasingly strict, leading to sharply improved air quality in many places. Research shows that better air quality improves neonatal health and reduces infant mortality (see Currie, 2009). But previous studies have not generally been able to assess the long-run effects of exposure to poor air quality in early life on the surviving infants and children.

Since relatively few children are on the life/death margin at birth, the total cost of air pollution could be much higher if other effects are taken into account. For example, many pollutants are 'neurotoxicants', which even at low levels of exposure can impair children's development. Psychological or behavioural problems not noticeable at birth or in early childhood may become apparent later on.

My study focuses on the causal impact of early childhood air lead exposure on later life outcomes (Nilsson, 2009). Merging unique data on local air lead levels in early childhood with comprehensive population micro data, it is possible to follow all the children in nine birth cohorts from birth through school and examine their experiences in the labour market as young adults.

The outcomes considered include scholastic performance, cognitive ability test scores, educational attainments and early labour market outcomes. These outcomes have previously been shown to be predictive of subsequent outcomes through the life cycle, and should therefore be particularly interesting from a public policy perspective.

Since the blood lead levels of Swedish children at their peak in the early 1970s were, on average, *lower* than the current limit suggested by the World Health Organisation, Sweden's experience is particularly interesting. It essentially provides a direct test of the relevance of the currently prescribed limit.

The local air lead pollution measure stems from a previously unexplored data source. Since the early 1970s, the Swedish environmental protection agency has conducted 'moss' surveys with a high spatial resolution across Sweden to examine regional differences and trends in heavy metal depositions. Moss lead levels are a good predictor of blood lead levels in children (Nilsson et al, 2009).

Although there is a consensus on the health impact of high levels of lead exposure on adult health, the association between lower levels of lead exposure in childhood and cognitive development is still under debate. The main reason is that lead exposure is not randomly distributed across locations.

Lead exposure early in life has significant effects on future educational attainments and labour market outcomes

Parents with high incomes or preferences for cleaner air are likely to sort into areas with better air quality and hence their children are less likely to be exposed to high levels of lead pollution. Failing to account for residential sorting of this kind can result in an overestimated effect of lead exposure on children's subsequent outcomes.

To mitigate this concern, my study focuses on children born between the early 1970s and the mid-1980s. The reason is that during the 1970s, along with many other developed countries, Sweden initiated a gradual phase-out of leaded gasoline to protect the environment and public health.

The main reduction in gasoline lead levels in Sweden occurred between 1973 and 1981, when the maximum lead level per litre of gasoline dropped by 79%. Since gasoline lead was the main source of lead exposure in the general population, children's blood lead levels decreased drastically from the 1970s until the mid-1990s when leaded gasoline was finally banned.

Due to substantial differences in initial lead levels, the phasing out of leaded gasoline meant that there was substantial variation across localities in the reduction of lead exposure. My analysis exploits the differential changes in early childhood lead

What are the long-run effects of exposure to poor air quality in early life? *Peter Nilsson* looks at the later life outcomes of people born in Sweden in the 1970s, a period in which regulations on leaded gasoline gradually reduced the lead content in the air and in children's blood.

exposure for the cohorts born between 1972 and 1984.

I compare changes in the outcomes of children born in municipalities that experienced large reductions in lead exposure with changes in the outcomes of children born in municipalities with only minor changes in air lead levels. By exploiting these differential changes in exposure across birth cohorts within the same municipalities, unobserved differences between the municipalities are taken into account.

My results suggest that low levels of lead exposure early in life have both statistically significant and economically important effects on future educational attainments and labour market outcomes. A key finding is a seemingly nonlinear relationship between local air lead levels in childhood and long-term outcomes at the relatively low levels of exposures considered.

Above an estimated municipality average early childhood blood lead level of 5 microgram/dL, reductions in lead exposure have a consistently positive and significant impact on long-term outcomes. Below this level, further reduction no longer seems to affect adult outcomes in a consistent or significant direction.

Children from poorer families have benefited relatively more from gasoline lead reductions

The estimated effects imply that by reducing average blood lead levels between ages 0-3 from 10 microgram to 5 microgram/dL, for example, high school graduation rates would increase by 2.3% and earnings at age 30 would increase by approx. 0.5-1% on average.

I also find that children from poorer families seem to have benefited relatively more from the gasoline lead reductions. Although data constraints prohibit a full differentiation of the mechanisms behind the differences in socio-economic status, a key finding is that residential segregation within municipalities (and thereby potentially differential neighbourhood lead exposure levels) *does not* seem to be able entirely to explain the different effects of lead.

Instead, different avoidance behaviour, differences in sensitivity to the same levels of exposure or differences in the ability to compensate for the effects of early lead exposure seem to be more plausible explanations for the relationship between socio-economic status and the impact of decreasing lead exposure.

Environmental policies may act as an instrument of redistribution by improving long-term outcomes disproportionally for disadvantaged children

Further analysis will potentially reveal which one of these pathways matters most. The current results indicate that environmental policies may be able to reduce the intergenerational correlation in economic outcomes. They could potentially function as a redistributive instrument since they seem to improve long-term outcomes disproportionally among children of lower socio-economic status.

Peter Nilsson is at the Uppsala Center for Labour Studies, Department of Economics, Uppsala University.

Further reading

Currie J (2009) 'Healthy, Wealthy, and Wise: Socioeconomic Status, Poor Health in Childhood, and Human Capital Development', *Journal of Economic Literature* 47(1): 87-122.

Nilsson, JP (2009) 'The Long-term Effects of Early Childhood Lead Exposure: Evidence from the Phase-out of Leaded Gasoline', mimeo, Uppsala University.

Nilsson, JP, S Skerfving, S Stroh and U Strömberg (2009) 'Pollution by Lead as Reflected in Moss and Children's Blood', mimeo, Uppsala University.

The effect of maternal fasting during pregnancy

Muslim women who are pregnant during the holy month of Ramadan will typically continue the practice of fasting during daylight hours. *Douglas Almond* and *Bhashkar Mazumder* investigate the impact on their children – in terms of birthweight, the likelihood of being a boy or girl, and later life health outcomes.

People's early childhood environment, including that experienced before birth, can have permanent effects on their lives. Recent studies by economists have used external shocks, such as exposure to famine or infectious disease, to provide compelling evidence in support of the 'foetal origins' hypothesis. An unresolved question is whether more commonly encountered early life exposures also have significant long-term effects on health and human capital.

In this study, we consider a common early childhood exposure: disruptions to the timing of pre-natal nutrition from skipping meals during pregnancy. Specifically, we consider the effects of maternal fasting during daylight hours in the month of Ramadan. Three in four Muslim pregnancies overlap with Ramadan, and surveys indicate that the majority of pregnant Muslims observe the fast.

There is currently no consensus on whether pre-natal fasting harms newborn health. We provide new evidence on the effects of fasting on birth outcomes and the first evidence of effects later in life using large sample micro data on Muslims in the United States, Iraq and Uganda.

Our methodological approach addresses a key flaw in previous studies of Ramadan fasting and birth outcomes. These epidemiological studies have compared pregnant women who fasted to those who did not at a point in time, under the basic assumption that the decision to fast is exogenous. Instead, we compare births over many years where Ramadan overlaps with pregnancy to those where Ramadan does not, and estimate the effect of Ramadan's timing.

This approach yields distinct estimates for specific months of gestation; Muslim births where Ramadan falls in the early post-natal period serve as the control group. Because Ramadan follows a lunar calendar, its occurence moves forward by roughly 11 days each year. Thus, we can disentangle the effect of pre-natal overlap with Ramadan from season of birth, which is also related to health in adulthood.

Pre-natal exposure to Ramadan results in lower birthweight, a reduced probability of a male birth and an increased likelihood of learning disabilities in adulthood

Using natality data from Michigan, we find that pre-natal exposure to Ramadan lowers birthweight and reduces gestation length. Furthermore, the likelihood of a male birth is about 10% lower when Ramadan falls very early in pregnancy and occurs during the peak period of daylight fasting hours. Using census data for the United States, Iraq and Uganda, we find long-term effects on adult health and some economic outcomes.

We generally find the largest effects on adults when Ramadan falls early in pregnancy. Rates of adult disability are roughly 20% higher, with specific mental disabilities showing substantially larger effects. Our estimates are conservative to the extent that Ramadan is not universally observed. Importantly, we detect no corresponding outcome differences when the same design is applied to non-Muslims.

Our findings are plausible in the context of biomedical research, where studies have documented that even relatively short fasts lasting 12 hours result in dramatic changes in the metabolic biochemical profiles of pregnant mothers.

This article summarises 'Health Capital and the Prenatal Environment: The Effect of Maternal Fasting During Pregnancy' by Douglas Almond and Bhashkar Mazumder.

Douglas Almond is at Columbia University. Bhashkar Mazumder is at the Federal Research Bank of Chicago.

CMPO working papers

Recent additions to the CMPO Working Paper series all of which, along with the rest of the series, can be downloaded from the CMPO website:

The Dynamics of Women's Labour Supply in Developing Countries, Sonia Bhalotra and Marcela Umaña-Aponte, 10/235

Housing Wealth or Economic Climate: Why Do House Prices Matter for Well-being? Anita Ratcliffe, 10/234

Is there an Income Gradient in Child Health? It Depends Whom You Ask, David Johnston, Carol Propper, Stephen Pudney and Michael Shields, 10/232

The Intergenerational Transmission of Worklessness in the UK, Lindsey Macmillan, 10/231

Intergenerational Persistence in Income and Social Class: The Impact of Within-group Inequality, Jo Blanden, Paul Gregg and Lindsey Macmillan, 10/230

Genetic Markers as Instrumental Variables: An Application to Child Fat Mass and Academic Achievement, Stephanie von Hinke Kessler Scholder, George Davey Smith, Debbie Lawlor, Carol Propper and Frank Windmeijer, 10/229

Digital Technology and the Allocation of Ownership in the Music Industry, Maija Halonen-Akatwijuka and Tobias Regner, 09/228

An Economic Theory of Glass Ceiling, Paul Grout, In-Uck Park and Silvia Sonderegger, 09/227

You Don't Always Get What You Pay For: Bonuses, Perceived Income, and Effort, Wendelin Schnedler, 09/226

Budgetary Separation of Powers in the American States and the Tax Level: A Regression Discontinuity Design, Leandro de Magalhães and Lucas Ferrero, 09/225

Parental Choice of Primary School in England: What 'Type' of School do Parents Choose?, Simon Burgess, Ellen Greaves, Anna Vignoles and Deborah Wilson, 09/224

Family Income and Education in the Next Generation: Exploring Income Gradients in Education for Current Cohorts of Youth, Paul Gregg and Lindsey Macmillan, 09/223

What Parents Want: School Preferences and School Choice, Simon Burgess, Ellen Greaves, Anna Vignoles and Deborah Wilson, 09/222

Test Scores, Subjective Assessment and Stereotyping of Ethnic Minorities, Simon Burgess and Ellen Greaves, 09/221

The Causal Effect of Education on Wages Revisited, Matt Dickson, 09/220

Do as the Neighbours Do: The Impact of Social Networks on Immigrant Employment, Fredrik Andersson, Simon Burgess and Julia Lane, 09/219

Gradients of the Intergenerational Transmission of Health in Developing Countries, Sonia Bhalotra and Samantha Rawlings, 09/218

The Centre for Market and Public Organisation
Bristol Institute of Public Affairs
University of Bristol
2 Priory Road
Bristol
BS8 1TX

Telephone: 0117 33 10799
Fax: 0117 33 10705
Email: cmpo-office@bristol.ac.uk
Website: http://www.bristol.ac.uk/cmpo/

CMPO | University of BRISTOL | E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL | The Leverhulme Trust