Choice-Based Secondary School Admissions in England: Social Stratification and the Distribution of Educational Outcomes

Rebecca Frances Allen
Institute of Education, University of London

PhD Thesis

November 2008

I hereby declare that, except where explicit attribution is made, the work presented in this thesis is entirely my own.

Word count (exclusive of appendices, list of references and bibliography):

73,205 words

Abstract

This thesis uses the cohort of 2005 school-leavers in the National Pupil Database to present an economic analysis of the effects of secondary school admissions in England on pupil sorting and achievement.

The first part of the thesis exploits the availability of pupil postcodes to examine the impact of current school admissions arrangements on residential and school stratification. It produces data from a thought experiment whereby pupils currently in schools are reallocated to a new school based strictly on proximity to school. Through this simulation the role of the housing market in producing socially stratified schooling can be identified. A survey of school admissions policies is used to show that religious (and to a lesser extent Foundation) schools have intakes that are significantly more advantaged than their local neighbourhood, and that they achieve these intakes through the use of explicit potentially selective admissions criteria.

The second part of the thesis adds to evidence on whether policies intended to increase parental choice raise pupil achievement via competition between schools. Quantitative evidence on school competition in England is evaluated alongside the existing international literature. A regression discontinuity design is employed to examine the legacy effects of the Grant-Maintained schools policy on area-wide educational outcomes at age 16. Pupil fixed-effects test score growth models and historical instrumental variables are used to identify the causal impact of religious schools on their neighbouring schools. This econometric analysis fails to lend support to claims that encouraging schools to compete for pupils is a route to improving standards.

The thesis concludes that the current English secondary school arrangements have resulted in a system that is stratified and inequitable, without measurable efficiency gains induced by competition between schools for pupils.

Acknowledgements

I wish to acknowledge the Economic and Social Research Council for financial support throughout my PhD. Thanks are also due to the Department for Children, Schools and Families for facilitating use of the National Pupil Database and Anne West at the Department of Social Policy, LSE, for providing access to her survey of school admissions policies.

I have been fortunate to work with many academics within and beyond the Institute of Education over the course of writing this thesis, and my discussions with them have added greatly to my ideas.

Early discussions with Rosalind Levačić about market reforms in education inspired me to pursue this research and I also thank her for reading a final draft of this thesis. Anne West has been very kind to me, particularly over my final year of study, and it would not have been possible to write Chapter 5 without her ideas and data. Simon Burgess, Frank Windmeijer and Harvey Goldstein have provided stimulating and ongoing collaborations in the measurement of segregation that have influenced methods used in Chapter 4. Rob French, Denise Hawkes, Ruth Lupton, and Jenny Neuburger all read chapters at various stages of development, and I am grateful for their comments. Graham Hobbs has been the perfect academic friend who is prepared to talk about school choice for hours each week, gently pointing out flaws in my thinking and analysis.

Special thanks are due to my supervisor, Anna Vignoles, who has influenced the development of this entire thesis, and has been a source of both professional and personal support throughout the last four years.

Finally, I would like to thank friends and family for their ongoing support and kindness. I could not have completed this thesis without them.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

> The debate about admissions, while often appearing to be about arcane technicalities, does in fact go to the heart of current policies about how best to achieve social justice, an improved education system and a cohesive society.
>
> (Coldron et al., 2008, page 3)

This thesis is concerned with the question of who benefits under the current system of secondary school admissions in England, where parents are given a choice over where their child is educated. The introduction of a quasi-market for school places has been central to government aims to improve school standards for the past twenty years. However, opinions remain polarised on whether these types of reforms are welfare enhancing. Critics argue that higher-income families benefit at the expense of the poor, because they are advantaged in their ability to exercise choice or because schools that control their admissions 'cream-skim' easier to teach pupils. By contrast, proponents suggest that it is the low income families who necessarily benefit most from the removal of the housing market's role as a gatekeeper to schools, because they are the group for whom '*the current constraints [of school allocation] are most binding*' (Hoxby, 2003a, page 10). Indeed, it is certainly untrue that selection by mortgage for neighbourhood schools is more equitable or less stratified than all conceivable choice regimes. Many advocates go further and argue that, regardless of how choice alters pupil sorting, all pupils ultimately benefit because competition between schools for pupils induces them to increase their effort, thus raising productivity (described by Hoxby, 2003b, as '*the tide that raises all boats*'). This thesis examines both the theoretical basis and the empirical evidence for these claims.

The chapters that follow focus on the sorting and achievement effects of secondary school choice in England. Choice is a process whereby parents are first asked to

express a preference for secondary schools, then school admissions criteria and practices determine how places for 11 year olds are allocated across the English state maintained sector, taking some account of parental preferences. Decades of cumulative reforms have resulted in a process that is quite complex for parents, with experiences varying considerably across the country. For example, in areas such as Cornwall, admissions to all maintained schools are determined and administered by the local authority, with priority essentially given on the basis of catchment areas. Therefore, choice of school is almost entirely related to choice of residential location. By contrast, in London 50 per cent of schools determine their own admissions, either because they are Voluntary-Aided schools with religious criteria for admission or because they opted out of local authority control during the early 1990s and are now Foundation schools.

The means and the ends of admission policies can be investigated in terms of many criteria, including school quality, efficiency, costs, responsiveness, accountability and equity (Le Grand, 2007). This thesis principally assesses the effect of secondary school admissions on equity, measured in terms of social stratification, and on efficiency, measured by average pupil achievement in GCSE exams at age 16.

Social stratification across schools can be seen as a direct measurement of levels of equality of opportunity in the system, defined as the extent to which a child's educational opportunities depend upon either the parents' economic circumstances or the child's location within the state (Wise, 1967). The secondary school a child attends is an important determinant of their GCSE performance[1] and the social class gap in achievement widens as children progress through secondary school. School allocation is one important determinant of these social class differences in educational progress.[2] Indeed, these inequalities in progress are an almost inevitable consequence of social stratification between schools if a child's peers affect their educational success.

Average GCSE achievement captures a relatively narrow aspect of efficiency since it ignores any possible effects of policies on educational costs and does not measure changes in the effect of schools on other pupil outcomes, such as learning not captured by exam papers, happiness, self-esteem and future wages. However, this focus on test scores is justified on the basis of its likely correlation with other outcomes: if children are happy and motivated in school, this should feed into superior GCSE achievement; if children are more successful in the labour market as a result of their schooling, this

[1] Reynolds (1992) report that the variance of achievement attributable to UK schools is 8-15 per cent. Wilson and Piebalga (2008) report that over 7 per cent of the variance in GCSE achievement is attributable to schools, even once peer group characteristics are accounted for, although 50 per cent of secondary schools are indistinguishable from each other in terms of quality on contextual value-added (CVA) rankings.

[2] Hobbs (2007) finds around one-fifth of class differences in Key Stage 2 progress are explained by class differences in school allocation. Sacker et al. (2002) find 25 per cent of class differences in achievement at age 11 were accounted for by class differences in school/classroom composition, but this estimate is likely biased upwards.

should have also been reflected in achievement at age 16.

1.1 Outline of thesis

The thesis is divided into two separate substantive questions:

1. What is the effect of school admissions policies on patterns of school and residential sorting?
2. What is the effect of these policies on the level of pupil achievement at the end of compulsory education?

These questions are distinct, yet closely related. However, questions of sorting and pupil achievement are addressed separately for the following reasons. First, it is possible to identify the effect of policies on sorting, without consideration of how the policies might influence school effort and pupil achievement. This first research question in itself relies on a considerable theoretical literature and methodological consideration of the measurement of sorting. By contrast, it is only possible to conceptualise how policies are likely to influence the level and distribution of pupil achievement once patterns of sorting are described. Thus, the ability to answer the second research question is contingent on addressing the first. Both research questions are important from a social welfare perspective: patterns of and reasons for changes in pupil sorting allow us to assess the equity of the schooling system; the level and distribution of pupil outcomes not only identifies whether the system is equitable, but also whether it is efficient.

Chapters 3 to 5 address the question of patterns of school and residential sorting in the current English schooling system. Chapter 3 provides a theoretical framework for predicting how pupils are likely to be sorted into schools under different admissions arrangements. It argues that the sociological post-structuralist literature has contributed greatly to our understanding of why choice strategies are likely to differ by social class, but is incapable of making comparative predictions as to levels of social stratification under alternative school admissions arrangements. By contrast, general equilibrium models of school choice are capable of making specific policy predictions, but the literature as it currently stands is highly reductionist and not well adapted to the English institutional environment. Chapters 4 and 5 provide descriptions of this English institutional environment, detailing how the housing market, school admissions criteria and parental strategies interact to produce current patterns of pupil sorting. Chapter 4 measures the contribution of the housing market to producing a stratified schooling system in England. It does so by using the availability of pupil postcodes to enact a 'thought experiment' that asks by how much school segregation could potentially reduced if a central planner intervened and allocated all pupils to secondary schools

strictly on the basis of proximity alone. Chapter 5 builds directly on the tentative conclusions of Chapter 4 by using a survey of school admissions policies to examine why some schools have intakes that differ substantially in terms of social and ability composition to the neighbourhood in which they are located.

Chapters 6 to 8 address the more broad and complex question of whether operating a quasi-market for school places raises pupil achievement. Understanding models and patterns of sorting is important to answering this question, for two reasons. First, the nature of demand for schooling determines how and why schools might compete for both quantity and type of pupils. Second, any changes in stratification as a result of quasi-market policies will affect which pupils benefit from the reforms. Chapter 6 describes a theoretical model of how school competition is likely to work in the English context, and it examines the existing international evidence for productivity gains resulting from school competition. Chapter 7 employs a regression discontinuity design to analyse whether the policy to allow schools to become Grant-Maintained in the early 1990s, and thus operate independently from local authority control, has been successful in increasing effort focussed on pupil achievement in the long-run, both through changes in the productivity of their successor Foundation schools and through competition effects on neighbouring schools. Chapter 8 examines whether the competitive threat that religious schools present to neighbouring schools raises, or lowers, area-wide academic standards. Pupil fixed-effect test score growth models are combined with historical instrumental variables to ensure the supply of religious school places in an area is independent of present-day demand.

1.2 Motivation for theoretical approach

This thesis draws on several theoretical perspectives, but is principally written in an explicitly economic and quantitative tradition. The origin of the economics of education as a significant field within economics dates back to the theoretical and empirical developments made by American economists such as Gary Becker, Jacob Mincer and Theodore Schultz in the 1960s (Machin, 2008). The crucial starting point of this economic analysis is that parents and school managers are rational agents (often given the name '*homo-economicus*'), operating with self-interested intent as they negotiate the school admissions process (Le Grand, 2003).

The market for secondary school places has been termed a 'quasi-market' because it differs in some fundamental aspects from classical free markets in respect to both the demand and supply sides (Glennerster, 1991; Le Grand, 1991). It is a market in the sense that parents act as consumers with a free choice of school. On the supply side, schools have budgetary independence and are encouraged to compete for stu-

dents through the funding formula. However, the state retains control of financing the service, thereby divorcing the distribution of the service concerned from the distribution of income (Le Grand, 2003). Furthermore, schools cannot easily expand or exit in response to consumer behaviour, do not earn profits and are subject to considerable operating restrictions via the National Curriculum and controls on teacher pay (Glennerster, 1991).

There are two further unusual aspects of this quasi-market. The first is that the nature of the demand for schools partially determines the quality of supply because peer characteristics (such as ability and behaviour) are an important input to the education production function. This means the choice of any one parent imposes an externality on all others since the allocation of their child to a school influences the overall distribution of peer effects. The second unusual aspect of this quasi-market is the absence of price as a rationing device for school places. This, combined with supply that is constrained for agency and institutional reasons, can produce a permanent mismatch between demand and supply for places at any one school. The school admissions policies at the centre of this study are simply the means by which places are currently rationed in England.

There are shortcomings of allowing parents to choose school places for their children, given the complexity and importance of the decisions to be taken. Most parents make infrequent, but closely-spaced, choices with the implications of these decisions only revealed in the long term, if at all (Levin, 1991). The decisions would optimally be made with an assessment of the probability of achievement at different schools, the relative costs of travel to schools, knowledge of school ethos or climate, the chances of successfully securing a place at each chosen school and evidence about labour market returns to education. However, it would be impossible for even the most well-informed parent to collect these data. The problem of imperfect information is then compounded by inadequacies in the ability of parents to compute information available to make the most appropriate choice for their child.

Despite these shortcomings of parents as decision-makers, there are theoretical economic models and empirical evidence that suggest that parents do a better job choosing schools for their children than experts in government agencies (Bast and Walberg, 2004). If we believe that parents are better able to choose the highest quality schools, this should result in the quasi-market for schooling being more allocatively efficient than the administrative allocation of pupils to schools. Schools should adapt their product to meet the desires of parents and popular schools would expand and less popular ones would shrink in size. However, if parents (correctly or incorrectly) think that peer groups form a significant part of their child's education production function, they necessarily care who their child goes to school with. Therefore, schools with

affluent intakes will be popular almost regardless of their effort.

This has two implications. First, parents may not choose schools that are the most productively efficient, meaning these schools do not maximise outcomes holding constant a set of inputs that *includes* the peer group. Thus, a process by which popular schools expand at the expense of others may be as likely to reduce productive efficiency as it might increase it. Second, it means that schools can best ensure they survive and prosper simply by concentrating their effort on securing an advantaged intake through the application of certain admissions policies and procedures (i.e. by cream-skimming). This, combined with evidence showing that family background appears to be important in schooling decisions, suggests that a quasi-market for school places risks becoming quite stratified.

1.3 Motivation for a quantitative approach

The empirical analysis that follows in Chapters 4, 5, 7 and 8 uses large scale administrative datasets containing information on all pupils in state maintained secondary schools to draw generalisations about the effect of policies across England as a whole. It is able to exploit the National Pupil Database (NPD), available from 2002 onwards, to reanalyse policy questions previously investigated only in aggregated schools data. Examples of this earlier work include the evaluation of changes in school segregation carried out by Stephen Gorard and others, and economic studies of school competition (e.g. Bradley et al., 2001; Gorard et al., 2003; Levačić, 2004).

Analysis of these research questions using pupil-level data brings several advantages. The analysis is likely to be more robust because the dataset provides variables measuring the characteristics of the pupils in the schools including, most importantly, the attainment of pupils prior to entering secondary school. Thus, the risk of confounding the effect of policy interventions by other factors is reduced. NPD also allows analysis to take place on a cohort, rather than whole school, basis so that the effect of changes in policies can be cleanly identified between two years. But aside from claims of greater validity of analysis, the pupil-level analysis also allows quite different questions to be addressed. Examples of this are the analyses in Chapters 4 and 5 that exploit the availability of pupil postcodes to provide a description of where pupils go to school, relative to where they live.

The research questions posed in this thesis are theory-testing rather than theory-building and so lend themselves directly to fixed design quantitative analysis, of which there are two distinct types in this thesis. The choice and sorting part of the thesis uses a descriptive or correlation approach to analyse school admissions, residential locations and school stratification. It seeks to identify statistical associations between

admissions regimes and pupil sorting. The analysis is not able, on its own, to make causal claims about the effects of policies on sorting since there is no robust identification strategy that is able to isolate the policy intervention from the context into which it was introduced. Where causal relationships are tentatively inferred in concluding remarks, this is only possible because the research questions are grounded by a combination of economic theory and existing smaller scale qualitative and quantitative studies.

The second part of the thesis uses quasi-experimental methods to make causal claims about the effect of school governance on competition and pupil achievement. It does so by using econometric techniques – a regression discontinuity design, pupil fixed-effects models and instrumental variables – to isolate and measure the effect of the policies on outcomes. This type of analysis is necessarily reductionist, assuming that all 'context' can be cancelled out via the use of a control group or variables. This position is challenged some who argue that a constant and controlled relationship between treatment and outcomes (known as the 'constant conjunction') cannot exist when dealing with humans and social situations in which a treatment is necessarily always applied in a unique context. Furthermore, though econometric techniques can successfully produce quantitative verification for theory, the reverse is not necessarily true where analysis yields equivocal results. The theory cannot automatically be rejected since the null result might be due the nature of the implementation of the policy, the context into which the policy was introduced or problems with the specification of the research design itself.

The key threats to the reliability, generalisability and validity of the analysis in this thesis are as follows. First, internal validity is threatened by complexity in operationalising key theoretical constructs, and in particular measuring levels of choice, competition and sorting in the system. Second, there is a risk to internal validity because all analysis is carried out on a single national cross-section of pupils, which clearly has major limitations when trying to describe a constantly evolving quasi-market. It means that between-area comparisons of policies take place, as opposed to exploring the evolution of one area over time. This makes it extremely difficult to remove context from associations between variables, and the control variables available may not prove sufficient to achieve this. Finally, even if we are able to claim validity to the target construct – all pupils passing through the English secondary education system around the start of the 21st century – the thesis is necessarily limited in its ability to extrapolate findings beyond this narrow context. All causal statements are inevitably contingent, but in education research the wide variation in persons, settings, treatments and outcomes across countries and over time mean that evidence for a theoretical position that is upheld in this study is of limited use in making policy predictions in other contexts.

1.4 Policy background

The final section of this introductory chapter gives an overview of the major policy events that led to the current English system of school governance and admissions. Mass secondary schooling in England is a 20th century phenomenon, and so the early history of schooling is not covered here. However, the role of the churches in the provision of 19th century elementary schooling is described in some detail in Chapter 8 since it is directly relevant to the estimation strategy employed in the analysis.

The current secondary school system, which children usually pass through from age 11 to age 16 or 18, has its roots in the 1944 Education Act. Prior to World War II, compulsory education lasted from age 5 through to 14. Slightly under half of all children completed this education in all-through elementary schools, with the rest transferring at age 11 to separate senior elementary schools/departments or grammar schools. All schools in the maintained sector – including church schools – were managed by one of 318 Local Education Authorities (LEAs), which were established as part of the 1902 Education Act to replace the Dual System of church provision and school boards. Only around 20 per cent of pupils had access to post-14 education via fee-charging day grammar schools, a small number of local authority secondary schools or through the entirely separate system of preparatory and public (private) schools (Barber, 1994). Few scholarships existed for those families who could not afford fees.

1.4.1 1944 Education Act

The 1944 Education Act heralded the start of a system of universal, free and compulsory secondary education, backed by widespread popular and political support. Although it took many years to be fully implemented by local authorities, it established the principle of free education in secondary schools, rather than elementary schools, until age 15. By 1951 there had already been a three-fold increased in numbers attending these secondary schools in England (Brown et al., 1997). The significance of the Act in providing an enduring institutional framework for the education service cannot be understated: most of its provisions remained actively in place until the recent legislative whirlwind that began in 1988.

A notable feature of post-war secondary provision was the decision to establish a Tripartite System of schools. A majority of pupils transferred to Secondary Modern schools, with a minority of pupils pursuing a more academic curriculum at grammar schools and a further (smaller) minority following a technical curriculum at Technical schools. Entry to grammar schools was free after World War II, but was conditional on passing the academic eleven-plus exam. It was recognised at the time of the 1944 Act that age 11 was probably too young an age to determine the future careers of children,

but the decision to make the transfer at 11 and not at 13 was largely made on practical grounds (Barber, 1994). Elementary buildings couldn't accommodate pupils up to the age of 13, and building work was infeasible, particularly in the religious sector where both the Anglicans and Catholics were in a state of financial crisis.

The Act gave LEAs considerable power to mandate diocesan education authorities, which had dominated the early system of elementary schooling, to reform their schools into primary and secondary schools. The negotiations with the churches are described further in Chapter 8, but an important feature was the creation of new governance status for church schools. They could become 'Voluntary Controlled' (VC), meaning that the State would take day-to-day control of the school. Alternatively, if the church wished to retain control of their school via 'Voluntary Aided' (VA) status, they were required to make a 50% capital contribution to the school. In return the church had majority representation on the governing body, giving them control over the employment of staff, buildings and repairs, and school admissions (Gay and Greenough, 2000). Thus, church and LEA schools continued to operate separately within the same system, although most church elementary schools became primary schools resulting in lower levels of church secondary schooling provision.

Parental choice of school was allowed, though restricted, under the 1944 Education Act. Section 76 stipulated that pupils should be educated in accordance with the wishes of their parents, but only in '*so far as is compatible with the provision of efficient instruction and training and the avoidance of unreasonable public expenditure*' (Ministry of Education, 1944). It has been argued that this was a relatively weak directive (Fitz et al., 2001). In any case, choice was understood at the time to implicitly signal denominational choice, rather than choice between secular schools since this clause was to appease concerns of the churches regarding access to denominational schools. Thus, choice as enshrined in the 1944 Education Act did allow a minority of parents with sufficient information and resources to satisfy their preferences within the state system, but the obligation to justify the selection of a non-local school still fell entirely upon parents (Walford, 1994).

1.4.2 Comprehensive schooling reform

The single most important reform to take place in the long period between the 1944 and 1988 Education Acts was the dismantling of the Tripartite system of secondary schooling. There was consistent opposition to grammar schools even at the time of the 1944 Education Act, but this became increasingly vocal in the early 1960s as the expanding middle class found their opportunities for a grammar school education increasingly competitive and unsatisfactory due to scarcity of places (Halsey et al., 1980).

Early comprehensive reforms took place through local initiatives, without a need for legislative mandates. Then in 1965, Circular 10/65 announced the Government's formal opposition to the eleven-plus exams and requested LEAs to reorganise on a comprehensive basis. Although the comprehensive ideal was emphasized in the Circular, it was neither put into a Green Paper nor legislation before the Conservatives regained power in 1970. At this time the political ground began to shift and Circular 10/65 was officially withdrawn when Margaret Thatcher became Secretary of State for Education. In 1976, the Labour Government passed an Education Act to finally introduce compulsory comprehensive re-organisation in LEAs. A second bill in 1978 required LEAs to use an admissions procedure based on 'parental preference', thus mirroring part of the spirit of the 1980s legislation, but it would never become law (Benn and Chitty, 1996). Although these Acts were reversed by Conservative bills after 1979, the new Government failed to prevent more local authorities from going comprehensive and phasing out selection even as late as the 1980s. Thus, the comprehensive sector had successfully grown from educating just one per cent of children in 1951 to 90 per cent in 1981.

1.4.3 Conservative government reforms – 1979 to 1997

Le Grand (2003) describes the English schooling system prior to the 1980s as a non-market trust model: local authorities exercised considerable influence in the allocation and spending of school funds and in the assignment of pupils to schools, yet teachers were given extensive freedoms to determine curricula and teaching methods. Critics of this model argued that government failure and agency problems surrounding teacher incentives meant that reform was needed. Pugh et al. (2006) describe the government failure argument as follows:

> Governments' desire to retain power leads to an emphasis upon short-term costs and benefits, time consistency problems and a preference for policies that have concentrated benefits and diffuse costs. Even in the absence of explicit corruption, public providers face incentives to favour government supporters and the interests of employees rather than cost reduction and innovation.
>
> (Pugh et al., 2006, page 20)

This non-market trust model of began to change over the course of the 1980s, starting with the 1980 Education Act and ending with the 1988 Education Reform Act. The 1980 Education Act made only modest changes of the right of parents to choose a secondary school for their child. It allowed parents to 'express a preference' for a school, but considerable LEA powers remained to manage falling school rolls and the

overall provision of places, taking into account benefits of the community as a whole (Department for Education, 1980). Nevertheless, they were required to make the means available for parents to express preferences for maintained schools, which shifted the burden of justification from parents to LEAs, who would now have to explicitly deny a preferred school if the parent requested one (Department for Education, 1992). However, choice was not perceived as existing in the many LEAs where parents were informed of their designated catchment area school first, with details of how to request an alternative place left as small print on the designation letter (Stillman and Maychell, 1986).

A more radical aspect of the 1980 Act was the introduction of the Assisted Places Scheme, which provided financial support to low-income parents to allow them to send their children to private schools. By 1987 almost 27,000 pupils were holding assisted places at an annual cost of some £48 million to the state and the scheme was not dismantled until Labour returned to power in 1997. However, very few of participants actually came from working class backgrounds: two-thirds of the mothers of the participants had themselves attended selective or independent schools (Edwards et al., 1989).

The quasi-market in its current form really emerged with the controversial 1988 Education Reform Act for England and Wales. The major thrust of the Act was designed to increase competition between schools and to encourage parents to make choices between schools. It was introduced against a backdrop of large demographic changes: by 1987 the number of ten-year olds had plummented by 30% from their peak in 1975, producing an abundance of surplus places in secondary school (Walford, 2003). This surplus capacity brought with it both the desire by parents to choose, and the ability for choice to be exercised in many areas. The 1988 Education Reform Act states that '*no child should be refused admission to a school unless it is genuinely full*', and locally assessed admissions numbers were introduced to ensure LEAs could not retain empty places at popular schools (Department for Education and Science, 1988). The Act came into effect for admissions to schools from 1990 onwards, ending the dominance of formal catchment areas and proximity criteria as the principal administrative devices for the allocation of pupils to schools. Choice was initially restricted by LEA boundaries, but the 1989 Greenwich Judgement removed the right of LEAs to deny access to pupils on the basis of these boundaries in cases of reasonable proximity.

A separate part of the 1988 Education Reform Act established a formal National Curriculum Assessment, with new tests to be taken at ages 7, 11 and 14. These tests resulted in nationally published data on schools that, along with the standard GCSE and A level 'league tables', contributed to a new transparency in terms of school performance in public examinations. During the same period, Local Management of

Schools funding reforms turned budgets over from LEAs to schools, with 80 per cent of funding determined solely by pupil numbers (Levačić, 1998). This gave schools a stronger incentive to attract and admit more students. It also had the effect of weakening the power of the local authorities and teacher unions traditionally allied with the Labour party (Walford, 1994).

A second major policy theme running through Conservative education reforms was the creation of new sets of schools operating outside LEA control. These schools were designed to act as beacons of excellence for local LEA schools with the aim of raising standards (Department for Education and Science, 1986). The first of these policies was the creation of City Technology Colleges in the 1986 (No. 2) Education Act. They were directly modelled on American Magnet Schools, but never really took hold in England, with just 14 being created in total.

The Grant-Maintained schools policy, explained in some detail in Chapter 7, was a far more successful exercise in school autonomy. The 1988 Education Reform Act enabled existing maintained schools to hold a parental vote to opt out of local authority control. Over the course of the 1990s, one-in-six secondary schools achieved Grant-Maintained status. One key consequence of independence from LEA control was that they could administer and change their admissions policies and just over one-in-ten choose to apply to the Secretary of State for Education to select as much as 30 per cent of their intake by ability.

The 1993 Act introduced a further policy to encourage diversity in school provision. Schools were able to apply for additional funding to acquire a Specialist status and, if successful, were also able to select up to 10 per cent of their intake based on pupil ability in subjects such as technology, music, drama, and sport. Attracted by the additional funds that Specialist status brought, this was a popular policy with schools, although most did not take advantage of the opportunity to introduce selection. However, these three policies in school autonomy, introduced alongside a set of legacy grammar schools and a large number of religious schools, meant that the quasi-market for school places in England, rightly or wrongly, became associated with schools choosing pupils as much as parents choosing schools.

1.4.4 Labour government reforms – 1997 onwards

The Labour Party had originally opposed many of the quasi-market reforms to education, yet in Government from 1997 they demonstrated strong support for parental choice and school autonomy, with few of the 1980s reforms being reversed. The 1998 School Standards and Framework Act enabled existing selection to remain, though it curbed the introduction of new selection policies (only specialist schools are allowed

to introduce selection by 'aptitude' for 10 per cent of pupils). The Assisted Places Scheme was abolished, as was Grant-Maintained status, but these schools did not return to LEA control, with most becoming Foundation schools (see Chapter 7 for details). Thus, a party committed to 'standards not structures' in education quickly moved towards introducing a series of structural reforms in the areas of diversity of provision and school admissions policies.

Diversity of provision

A series of reforms over the past decade have focussed on changes in school type and governance away from the traditional Community school model. City Technology Colleges, Specialist schools and now, Academies, all form part of this current policy trend towards diversity of secondary school provision. The policy of diversity is seen to simultaneously achieve two policy objectives: it is consistent with the right of parents to educate their children according to their views and beliefs and it is a pre-requisite of an effective education market where genuine diverse choice facilitates competition between schools.

These policies – particularly the Specialist schools program – have altered schooling in every local authority in England. However, the most radical changes have been deliberately focussed on urban areas where schools were perceived to be failing. Most Academies – publicly funded independent schools under private or voluntary management that are allowed to select pupils – have been opened in disadvantaged urban areas. In doing so, they have been introduced against the policy backdrop of Excellence in Cities (funding targeted at pupils in schools located in disadvantaged areas) and Education Action Zones (which allow some schools to work with business sponsors). However, West and Pennell (2002) argue these programs have done relatively little to alter the urban schooling structure generated by the 1988 Education Reform Act.

Experience in secondary schools for the pupils under analysis in this thesis is still highly dependent on region and local authority. Almost all – 93 per cent – began their secondary school careers in year 7 at age 11 (and the number of middle schools has continued to decline since 2000) and most will continue in the same school until age 16 or 18. Two-thirds will be educated in a Community (or Voluntary Controlled) school that is controlled by the local authority. About 15 per cent will be in Voluntary Aided (VA) schools that are mostly owned by churches or charities and control staffing, capital spending and pupil admissions. Two-thirds of these are Roman Catholic and the majority of the rest are Church of England. Around 16 per cent will be admitted to Foundation schools, most of which are former Grant-Maintained schools. These schools are usually officially owned by the Governing Body of the school, and again control

staffing, capital spending and pupil admissions. Very small numbers of pupils are in technically 'independent' state-funded City Technology Colleges or Academies, where the local authority has no representation or monitoring function over the school.

School admissions codes

The Labour Government has introduced a series of legislative reforms to tighten the rules regarding how schools and local authorities may allocate their places to pupils each year. The 1998 School Standards and Framework Act established a new legal framework of regulation for admissions along with a revised Code of Practice for September 2000 admissions, designed to alleviate some of the problems arising from a largely unregulated education market (Department for Education and Employment, 1999). The new measures created included a new Adjudicator position for resolving local admissions disputes, as well as a prohibition on any new selection by ability in schools other than by a process of ability banding (DfEE, 1999). The 1999 Code of Practice stipulated that secular schools should not interview parents as part of their admission process, and that religious schools should interview only to establish an applicant's religious denomination and commitment, but there continued to be concern that this right to interview was being used for the purposes of social selection. It was under this admissions framework that the pupils in this study were admitted to secondary school.

Concerns over the fairness of school admissions continued and so in 2003 a new Code of Practice (Department for Education and Skills, 2003) was introduced with three mutually reinforcing measures intended to regulate admissions authorities:

- the requirement to consult;
- the requirement to have *regard* to the published guidance; and
- the opportunity to object to the Schools Adjudicator.

The 2003 Code of Practice continued to give schools considerable discretion in deciding their own over-subscription criteria, provided that they were lawful, impartial, publicly known and geared towards providing all local children with a suitable education. Schools were obliged to provide information and guidance for parents on admissions procedures and oversubscription criteria. Schools and local authorities were not allowed to apply criteria that had the effect of disadvantaging certain social groups. An adjudication process exists as part of the code to resolve disputes and there are rights of appeal to a Schools Adjudicator.

However, this Code was not legally binding on schools and so its success was mixed. For example, the 2003 Code extended the interview ban to include children, except for

boarding schools, but was unsuccessful in preventing the London Oratory (an inner city Catholic school) from holding pre-admissions interviews with prospective pupils. The quality of information provided by schools also failed to meet the requirement of the Code. So, although schools should report the '*number of places and applications for those places in the previous year*', Coldron et al. (2008) report that is was not possible to tell if the school was over- or undersubscribed at 42 per cent of schools.

In February 2007, the new School Admissions Code came into force, with even greater powers than either previous Code (Department for Education and Skills, 2007b). It introduced a number of strong measures that admission authorities must, or should, follow. Its requirements on admission authorities concerning such details as permissible oversubscription criteria, the prohibition of giving higher priority to those who place a school higher on their application form (usually referred to as first preference first) and requirements to publish certain admissions information should significantly impact on practice (Coldron et al., 2008).

The role of local authorities (LEAs were re-labelled LAs as part of the 2007 legislation) as co-ordinators has also simplified school admissions and prevented parents from holding offers for places from more than one admissions authority, thus denying other pupils places. They act as admissions co-ordinators who invite parents to express a preference for a minimum number of 3 schools (with 27 per cent of LAs inviting 6 or 7 preferences). This information is used in one of two ways. For the 101 authorities that use an equal preferences ranking system, the LA (or a pan-London body for the capital) asks all admissions authorities in the area to rank all parents who have expressed a preference for their school in order of priority based on published oversubscription criteria. A two-sided matching model known as the Gale-Shapley algorithm is employed to reconcile the preferences of parents for schools with the 'preferences' of schools for parents (Gale and Shapley, 1962).[3] This is a complex procedure whereby parents are given the best place possible for them, given the expressed preferences of others, in a series of rounds. In the 47 authorities that still used a first preference first system (now ruled out under the 2007 School Admissions Code), the authority asks schools to only consider the first stated preference of the parents, and to firmly allocate places where possible. Second preferences of parents are only considered where the first preference cannot be met. This makes it imperative that the parent acts strategically to ensure that they are highly likely to gain a place at their first preference school, since if they do not, their second and third preference schools are likely to be already filled by other parents who listed the school in first place.

Current knowledge about secondary school admissions is largely based on two surveys of school admissions criteria and practices. Anne West and colleagues collected survey

[3]There is a large economics literature on the stability and efficiency of two-sided matching mechanisms (see, for example, Roth, 1984).

data on over 95 per cent of comprehensive schools for the 2001 cohort entry (West et al., 2004). John Coldron and his team documented 2006 admissions for all state maintained secondary schools (except CTCs and academies) (Coldron et al., 2008). From these surveys it is known that twice as many schools now determine their own admissions as did at the time of Education Reform Act: 30 per cent compared to 15 per cent in 1988 (West and Hind, 2003). Nevertheless, catchments or distance to school criteria have roles of continuing significance in much of the country (Jowett, 1995; Mayet, 1996). Typical oversubscription criteria cited by schools in the study of all comprehensive schools by West et al. (2004) include siblings (mentioned by 96% of schools), distance (86%), medical/social needs (73%), and catchment areas (61%).

Grammar schools remain in 43 LAs, 20 of which have over 10 per cent of places available on the basis of ability selection. In addition, the 1998 Standards and Framework Act allows schools to continue with pre-existing selection arrangements. Coldron et al. (2008) found 10 'comprehensive' schools that selected more than 10 per cent of their intake. The number of school selecting 10 per cent of their intake on the basis of 'aptitude' for a subject has increased from 1.3 per cent in 2000, to 3 per cent in the 2001 admissions survey, to 4 per cent in a study of 2006 admissions. These partially selective schools are mostly VA or Foundation schools with a Specialist status.

Aside from explicit ability selection, opportunities for more covert social selection did exist in the 2001 and 2006 survey (and may still do so). They are described in more detail in Chapter 5. For example, own-admission schools were far less likely to give priority to children in public care, or children with medical or social needs for 2006 entry (Coldron et al., 2008). Four per cent of schools used criteria in 2006 relating to 'associated adults' and 9 per cent refer to proof of parental commitment in such a manner that would also now be unlawful.

Coldron et al. (2008) report that 29 per cent of non-Community schools were still using supplementary forms in 2006. Pennell et al. (2007, page 1), in their survey of London secondary schools, concluded that '*the supplementary forms used by some schools provide opportunities to select more desirable pupils*', with schools asking parents to provide information that did not seem to be related to the school's admissions critieria (e.g. whether parents were living in bed and breakfast accommodation or parents' occupations). They were also concerned that the length of some forms and the requirement, in some cases, that parents and children write extensively about their reasons for wanting a place at the school, could systematically deter some parents or carers from applying and such information could only be judged subjectively. The 2007 School Admissions Code now prohibits the use of supplementary forms which ask any personal information that is not relevant to applying acceptable oversubscription criteria.

Having documented these continuing concerns about the fairness of school admissions, it is easy to overstate the extent to which the current system of secondary school admissions is in a state of flux and crisis. Coldron et al. (2008) report that 85 per cent of parents gained their first expressed preference of secondary school in 2006, a similar figure to 2000. In addition, just 5 per cent (8 per cent in the Flatley et al., 2001, survey) of parents surveyed reported that there was another school they would have liked their child to attend, but did not even apply to. Of course, patterns of satisfaction with the current arrangements vary across the country and it is no surprise that the percentage expressing satisfaction is lower at around 70 per cent in London. It is generally lower in more urban areas and in areas with greater numbers of VA or Foundation schools or Academies (Coldron et al., 2008). Where parents reported dissatisfaction, this generally focussed on the outcome of the admissions process, rather than perceived unfairness in the process itself. This apparently high level of satisfaction with current arrangements contrasts with a report in 1996, which highlighted considerable difficulties in the management of the process of choice, with many parents feeling dissatisfied with both the process and the outcomes (Audit Commission, 1996). It is not clear whether the system has actually significantly improved, whether parental expectations have adapted, or whether the sampling strategy and questioning is responsible for these apparent rises in parental satisfaction with school choice.

Chapter 2

Sources of data

This chapter provides details of the data sources used throughout the thesis. A cohort of 2005 school leavers, taken from the National Pupil Database (NPD), is central to all empirical analysis in these chapters. It is supplemented by other datasets for three main purposes:

1. Small area statistics are attached to each child's postcode or output area as indicators of the type of street or area the child lives in. These are intended to proxy for the socio-economic circumstances of the pupil's family, given the lack of good variables for pupils' social background in the National Pupil Database.
2. Variables describing the historical and current characteristics of the local authority or county the pupil lives in are used as descriptive statistics and as control variables in regression analysis.
3. Surveys detailing school governance and admissions policies supplement the main dataset as explanatory variables of interest.

The remainder of the chapter describes each data source in turn. These are summarised in Table 2.1.

2.1 National Pupil Database

The main source of data in this thesis is the Department for Children, Schools and Families' (DCSF) National Pupil Database. This administrative dataset combines attainment data for pupils in Key Stage tests at ages 7, 11, 14 and 16 with a limited range of pupil characteristics, collected from maintained schools via an electronic return that includes named pupil records in January each year (since 2002, and termly from 2007 onwards).

Table 2.1: Summary of data sources

Source	Unit of observation
DCSF National Pupil Database (2005 school leavers, linked to 2003 KS3 tests and 2000 KS2 tests in primary school)	Pupil
Indices of Deprivation 2004	Output area
ACORN^{TM} indicators of household type	Output area
Office for National Statistics Census of Population 2001	Output area
DCSF Annual Schools Census (1999 to 2005)	Secondary school
DCSF Register of Educational Establishments (EduBase) (2004, 2005)	Secondary school
Grant-Maintained Schools Database	Secondary school
2001 survey of secondary school admissions policies	Secondary school
2005 English Church Census	Local authority
1851 Census of Religious Worship	Ancient counties
1851-1961 Survey of religious affiliations	Ancient counties

It is a statutory requirement on schools to submit this information under section 537A of The Education Act 1996 since it forms the basis for allocating funding to local authorities and schools. Its status as a statutory process means there are not significant ethical issues surrounding parental consent for collection of the data since it is not sought by schools. However, DCSF makes modifications to the database to ensure that researchers cannot identify individual children. This includes removing full names, recoding the Unique Pupil Number with a Pupil Matching Reference, replacing date of birth with month of birth, and restricting access to the pupil postcode field.

Pupil data are collected to identify gender, date of birth (month of birth is used), name, ethnic group, first language, special educational needs (SEN) categorisation, eligibility for free school meals (FSM), home postcode and status as a boarder or part-timer. Sweeps from years 8, 9 and 11 are used in this thesis. Key statistics for the main variables from the year 11 sweep are shown in Table 2.2

The pupil postcode can be used to place each pupil's home address on an Ordnance Survey (OS) grid location to within 1 metre of the mean postcode position and within 100 metres of the pupil's home address. This provides a relatively precise measure of where the pupil lives, relative to where they attend school.

Pupil postcodes also enabled small area data to be attached to the child's street or immediate neighbourhood, as described in the next section. The All Fields Postcode Directory provides codes for statistical areas which are used to merge these data (Office for National Statistics, 2004). The smallest statistical area is called an output area (OA), which contains approximately 150 households (an average of less than 4 cohort pupils). OAs are non-overlapping geographical areas which are generated by a

Table 2.2: Key pupil variables in NPD

Variable	Number of pupils	Proportion
Female	573,227	49.6%
FSM	570,949	13.3%
SEN statement	570,828	2.5%
SEN action	570,828	9.3%
SEN action plus	570,828	4.1%
English not mother tongue	569,841	7.4%
Ethnicity white British	570,951	83.3%
Ethnicity asian Indian	570,951	2.3%
Ethnicity asian Pakistani	570,951	2.3%
Ethnicity asian Bangladeshi	570,951	0.9%
Ethnicity black African	570,951	1.5%
Ethnicity black Caribbean	570,951	1.5%

14 ethnicity indicators in total and 14 age indicators

computer so that, as far as possible, they are wholly contained within the 2002 ward boundaries and minimise geographical area covered. OAs are combined (typically five) to create 'lower-layer super output areas' (Lower SOA) with a mean population figure of 1500 (an average of 17 cohort pupils). These can be again be combined to create 'middle-layer super output areas' (Middle SOA) containing an average of 78 cohort pupils. These statistical areas are more appropriate to use for analysis than wards, districts or counties because their boundaries are permanently fixed and they contain similar sized populations.

In theory, inaccuracies and missing data should not be a major issue in the dataset. However, they arise for four main reasons. First, although it is a duty of schools to ensure they find out the Unique Pupil Number for new children joining the school, in practice some schools use newly issued or temporary numbers that make matching prior attainment data impossible. Second, individual fields such as the pupil postcode or child's ethnicity can be missing, and again this is most likely in children who have recently moved schools. Third, Key Stage attainment data will be missing where a child did not sit a particular test, as discussed below. Fourth, Key Stage attainment data will also missing for children who have spent time out of the state maintained system because they were abroad or in private schools.

These administrative datasets are usually described as populations, rather than samples, meaning that no discussion of sampling technique and validity is needed. However, large groups of pupils are either missing or dropped from the analysis. Most importantly, data on private school pupils is not used in any analysis in the thesis. These pupils do appear in NPD on each occasion that the pupil takes a nationally administered test (i.e. KS4, KS5 and other Key Stages for a sub-set of private schools).

However, private schools do not submit pupil-level records as part of the annual census of schools, so NPD lacks almost all pupil background indicators. This is a clearly an issue in the analysis of school choice since parents who have exercised the choice to attend a private school are ignored, even though they may have actively sought, and failed to gain, a place at a particular state school.

Other pupils are included in NPD, but are excluded from the analysis chapters. Pupils in non-standard secondary education (special schools, hospital schools, detention centres) are never included in analysis. In addition, the analysis of sorting of pupils into schools excludes the very small number of pupils who are registered as borders and pupils on the Isles of Scilly and Wight where no pupils make daily journeys to the mainland.

2.1.1 Key Stage test data

Pupils in maintained schools are statutorily assessed, using national tests, at the end of three Key Stages (KS) at ages 7, 11 and 14. They are awarded a level on a scale of 1 to 8 plus 'exceptional performance' on each occasion, with expected performance rising from levels 1 to 3 at KS1, levels 2 to 5 at KS2, and levels 3 to 8 at KS3. Analysis in this thesis uses data from KS2 and KS3 (matched KS1 data was not available for this cohort). In addition, all pupils in the dataset have taken national GCSE tests at age 16 (also known as KS4).

The expectation is that all pupils will take all Key Stage tests, but approximately 10% of pupils have some missing data. A minority of pupils are exempt from a test because they are working at a level below that assessed by the tests, or they are disapplied (exempt from the National Curriculum under Section 364 or 365 of the 1996 Education Act). A larger group of pupils will miss one or more subject tests. Rather than lose these pupils from the analysis, their missing scores are imputed where possible, using a single inputation method, using test score information from the other test papers that they sat in the same subject. A further group of pupils will have taken the tests but will not have been assigned a level because they only achieved a very low score. However, their mark on the paper is available and so an approximately equivalent level is imputed.

The prior attainment of pupils on entry to secondary schools is measured using Key Stage 2 (KS2) test scores in maths, English and science. These are sat in primary school at age 11, so scores are independent of secondary school effectiveness. Information on levels and marks achieved by pupils are available, so the mark data are re-calibrated as fractional equivalent of levels to produce three continuous prior attainment variables. This re-calibration takes place by identifying the lowest mark for

each level and applying a linear recalculation of level for all pupils within each level. For example, if the lowest test scores for level 4 and 5 are 45 and 65 marks respectively, a pupil with 55 marks will be assigned a level 4.5.

Key Stage 3 (KS3) test data are also available in maths, English and science, and a similar process is used to re-calibrate the marks data into fractional equivalent levels. This process is slightly more complicated as pupils are entered into different tiered papers depending on the expected achievement. The range of marks at each level therefore varies for each of the tiered papers.

Pupil achievement in GCSE exams (KS4) is summarised by the total points achieved by pupils over their best 8 GCSE/GNVQ or equivalent examinations. The GCSE exam is graded from A* (58 points) to G (16 points), with 6 point increments between grades. The typical pupil achieves around 4 Cs and 4 Ds. In addition, grades achieved in maths, English (language) and science (best science score across any exam) are used. Data for all these Key Stage tests is summarised in Table 2.3. These pupils sat KS2 in the year 2000, KS3 in 2003 and GCSEs in 2005.

Table 2.3: Key pupil test score variables

Variable	N	Mean	Std. Dev.	Min.	Max.
KS2 English score	552,375	4.374	0.861	0.000	6.000
KS2 maths score	551,899	4.377	0.878	0.000	6.000
KS2 science score	551,904	4.662	0.701	0.083	6.000
KS2 total z-score	552,213	0.000	1.000	−6.011	2.055
KS3 English score	560,721	5.401	1.254	0.000	8.000
KS3 maths score	562,018	5.858	1.342	0.000	8.981
KS3 science score	561,981	5.512	1.088	0.000	8.000
KS3 total z-score	559,111	0.000	1.000	−4.641	2.405
GCSE English score	573,227	4.543	1.870	0.000	8.000
GCSE maths score	573,227	4.288	1.966	0.000	8.000
GCSE science score	573,227	4.218	2.016	0.000	8.000
GCSE capped score	573,227	290.604	105.379	0.000	540.000
GCSE capped z-score	573,227	0.000	1.000	−2.758	2.367

2.2 Indicators of pupil family background

The socio-economic circumstances of the child's family – including household income, parental education and occupation – is central to the analysis in this thesis since it is a key determinant of child's educational achievement and degree of choice of secondary school. This section describes the relative merits of the alternative measures of family

background used in the thesis.

2.2.1 Free school meals eligibility

The main drawback of using NPD is that it does not provide a good indicator of the socio-economic status of the child: the only available metric within the dataset is the child's eligibility for free school meals (FSM), an indicator of low income. Children are classified as 'eligible' for FSM only if they are both eligible for and claiming FSM. It is intended to proxy for low income since eligibility is contingent upon the family receiving other benefits/tax credits, such as Income Support, Income Based-Job Seeker's Allowance and Child Tax Credits.

However, according to Hobbs and Vignoles (2007), there is a poor relationship between FSM 'eligibility' and net household income. They find that only 23-55% of the 16% of children 'eligible' for FSM are in the bottom 16% of the distribution of household income in 2004/5. This is principally because the receipt of means-tested benefits (and tax credits) pushes children eligible for FSM up the income distribution. There is also a concern that the pupil's probability of applying for FSM eligibility status depends, in part, on the culture of the school (Croxford, 2000).

2.2.2 Indices of Deprivation and ACORN housing types

Two deprivation indices are attached to the Lower SOA for each pupil to proxy for the child's socio-economic circumstances. These are important additional indicators of SES, given that FSM is an imperfect measure of income and cannot distinguish between the large majority of pupils who are not FSM eligible.

The two continuous measures of a deprivation used are the Index of Deprivation 2004 and the Income Deprivation Affecting Children Index (IDACI) (The Office of the Deputy Prime Minister, 2004). They were produced by the Social Disadvantage Research Centre at the University of Oxford to update the Indices of Multiple Deprivation 2000.

The Index of Deprivation 2004 is a Super Output Area (Lower SOA) level measure of multiple deprivation and is made up of seven Domain Indices which relate to:

- income deprivation;
- employment deprivation;
- health deprivation and disability;
- education, skills and training deprivation;

- barriers to housing and services;
- living environment deprivation; and
- crime.

The IDACI measure of income deprivation affecting children is a sub-domain of the education, skills and training deprivation domain. Both these indices are imperfect indicators of the child's social background to the extent that they measure average social characteristics of households in the SOA. There is a specific problem with this in the context of choice and sorting analysis, which is why they are not used in Chapters 4 and 5. For example, the presence of a faith school in a town might allow a church-going family to purchase a less expensive house in the catchment area for a Community school with a deprived intake, knowing that their children would attend the faith school. Geo-demographic data are therefore likely to systematically underestimate the socio-economic characteristics of this family, thus biasing all parameters of interest in sorting models.

In addition to the deprivation indices, 57 dummy variables indicating household type are included as control variables in pupil achievement models. These ACORN (acronym for A Classification Of Residential Neighborhoods) dummies were devised by CACI Limited as a means of classifying areas according to various Census characteristics. These are attached directly to the child's postcode, thus being more disaggregated than the deprivation indices. They have been shown to explain a significant proportion of the variance in pupil achievement at GCSE (Webber and Butler, 2007). However, because they are produced by an organisation who will not describe their construction for commercial reasons, there has been little independent academic evaluation of their validity or appropriateness for use in analysis.

2.2.3 2001 Census of Population

Both the deprivation indices and the ACORN indicators are principally variables derived from data collected as part of the 2001 Census of Population. In addition, local authority level data are extracted directly from census returns to facilitate additional analysis. In Chapter 4 the population density of the area, the proportion of families that are lone parent households and the proportion of families where no parent is employed are used. In Chapter 8 the proportion of the population reporting that they are Christian, Hindu, Muslim, Jewish, Buddhist and Sikh are used as control variables because other religious sources are restricted to Christian denominations.

2.3 Annual Schools Census and Register of Educational Establishments

The secondary schools in this analysis are named in the Register of Educational Establishments and are classified in terms of their school governance, religious affiliation, selective policy (indicating the 164 grammar schools), and sex of intake. The Annual Schools Census also provides information on school size, aggregate pupil characteristics and basic staffing information from which the pupil-teacher ratio can be calculated. The school postcode is used to give each school easting and northing OS grid references. School governance and religious affiliation are summarised in Table 2.4.

Table 2.4: Governance of religious schools

	Non-religious	Religious	Total
Academy	10	4	14
CTCs	14	0	14
Community	1961	0	1961
Foundation	492	9	501
Voluntary-Aided	42	480	522
Voluntary-Controlled	55	38	93
Total	2574	531	3105

Academies and City Technology Colleges are not specifically analysed in this thesis since they are few in number and the Academies program was very new at the time these pupils started secondary school. Voluntary Controlled secondary schools are also relatively few in number, and are often grouped with Community schools in analysis since they do not control their own admissions.

A detailed description of the institutional differences between Community schools, Foundation schools and Voluntary-Aided schools can be found in Chapter 7 (Foundation schools) and Chapter 8 (Voluntary-Aided schools).

2.4 Grant-Maintained Schools Database

The Grant-Maintained Schools Database is used in Chapter 7 to extract voting data from the 1990s for all schools who held a parental ballot to opt-out of local authority control. The database was compiled by the Department for Education and is now archived in the National Digital Archive of Datasets. It provides details of all GM status ballots taken by schools, and the outcomes of these ballots. It also gives details of major changes to the status of GM schools, such as requests to change admissions policy or introduce a sixth form.

Table 2.5 shows that around one quarter of secondary schools in existence today held a parental ballot in the 1990s. One-in-six achieved independence from local authority control, including just over 10% of schools which had lost their first parental ballot to opt out and held a second or third re-ballot.

Table 2.5: Schools in Grant-Maintained schools database

	Won vote first time	Lost first vote	Never took vote
Number of pupils	111,840	47,541	413,846
Number of schools	618	233	2,257
Won subsequent re-ballot		10.8%	

Table 2.6 shows the current governance of former Grant-Maintained schools. The majority were designated with the new Foundation status, but those who were formally Voluntary-Aided status were able to revert back to this status.

Table 2.6: Governance of secondary school GM vote winners and losers

	Final GM ballot won	Final GM ballot lost	GM ballot never taken
Grammar	94	8	70
CTC/Academy	0	0	14
Community (non-grammar)	12	162	1750
Foundation (non-grammar)	404	1	22
Voluntary-aided (non-grammar)	133	26	331
Voluntary-controlled (non-grammar)	0	11	69
Total	643	208	2256

2.5 West 2001 school admissions survey

A survey of school admissions criteria used by state maintained secondary schools for pupils entering year 11 in September 2001 was carried out by Anne West and others at LSE. A database was created with details of admissions criteria used by 95 per cent of comprehensive secondary schools in England and was made available in electronic format. LEA and school admissions brochures were used as the principle source of information for the study. The pupils analysed in this thesis actually began secondary school one year earlier, in September 2000. However, there were no major changes in admissions policies across England and the website of the Schools Adjudicator confirms that no schools applied for major changes to their admissions during this period.

Full details of the survey and its main findings can be found in West et al. (2004).

In this thesis these data are principally used in Chapter 5 to look at the association between school admissions criteria and pupil sorting. The dataset is also used to supplement incomplete data on levels of selection in former Grant-Maintained schools, as documented in the Grant-Maintained schools database.

2.6 Church attendance data

This section describes the three sources of information regarding the historical and current religious composition of the population in England. These data are principally used in Chapter 8 as explanatory variables to predict the supply of religious schools across England, and also as control variables in analysis.

2.6.1 1851 Census of Religious Worship for Great Britain

The 1851 Census of Religious Worship for Great Britain was taken alongside the population census and an educational census on the night of Mothering Sunday (30th March). It was the first census of its kind to ask British churches of all denominations to report their number of places of worship, sittings (space in church) and levels of attendance at worship that Sunday. Data from this census are used in Chapter 8 to describe the historical sizes of the Anglican, Catholic and non-conformists denominations at a time when mass education was being established.

Data from the Religious Census were summarised by registration district, of which there were 4,542 in Great Britain, and deposited electronically on the National Digital Archive of Data by Paul Ell as part of his doctoral work whilst at the University of Leicester and University of Birmingham. This data deposit provides users with the uncorrected figures, as originally published by Mann in the Census of Great Britain 1851 (1960).

There are serious concerns regarding the accuracy of the census, which partially explains why the process was not repeated until recently (Pickering, 1967). The data are subject to serious measurement error because some churches failed to return surveys (and some returned more than one) and attendance is a poor measure of size of the congregation (in non-conformists congregations it was particularly common for members to attend twice on Sunday, thus inflating the congregation size statistics). Nevertheless, despite these data quality issues it is able to provide a rough description of the relative size of religious groups across Great Britain in the mid-19th century.

In Chapter 8, the dataset is aggregated to ancient county level and the relative sizes of the denominations are described across areas. Although there are adjusted statistics

of congregation size available from historians, the uncorrected figures as calculated by Ell directly from the Census are used.

2.6.2 1851 to 1961 survey of religious affiliation

The main source of historical religious population information is a survey of the UK between 1851 and 1966, collated by M. Hechter of the University of Washington and deposited in the UK Data Archive at the University of Essex (Hechter, 1976). The purpose of the original survey was to collect social, demographic, electoral and linguistic data for each of the 118 British and Irish counties (now known as ancient counties). His sources included data from general elections, population censuses and other published survey sources.[1]

Variables used in Chapter 8 describe the proportion of the ancient county population who are Church of England, non-conformist or Roman Catholic. These are reported once every 10 years, from 1851 to 1961 (excluding 1941). These data are used as background to describe the growth and decline of church attendance across regions and denominations. In addition, the variable indicating levels of Catholic affiliation in 1931 is used as an instrument to predict the modern-day distribution of Catholic secondary schools in England.

2.6.3 2005 English Church Census

The most recent English Church Census was carried out by Christian Research (with support from the University of Manchester) in 2005. It is a survey of church attendance over the weekend of the $7^{th}/8^{th}$ May across all Christian denominations. Church leaders were asked to fill in questionnaires detailing type of Christian denomination and mission, the size, age and ethnicity profile of the congregation present, and the types of activities the church is currently engaged in.

These modern English Church Censuses from 1979 onwards have presented an overall picture of quite dramatic decline in church attendance in England, with just 6% of the population estimated to have been at church over the weekend in question in 2005. The data are summarised by local authority in Brierley (2006), and is used in the thesis as descriptive and control variables in Chapter 8.

[1]The quality of these data was evaluated by Derek Unwin of the University of Bergen in 1977 and his comments are available in the UK Data Archive. He notes that the religious data should be 'treated with caution' because there appears to be significant variation in church attendance levels over time.

Chapter 3

Models of parental choice of school

The next three chapters provide explanations as to why English secondary schools are socially stratified, with children from socially disadvantaged backgrounds unevenly distributed across schools. There are two distinct potential causes of stratification arising from the supply of, and demand for, school places:

1. schools use overt admissions policies, such as selection by ability or religious adherence, or more covert procedures to 'cream-skim' more able or easier to teach pupils;
2. the actions of parents from different social backgrounds somehow lead to differences in how school choice and house location decisions are made.

Investigating the former suggestion that schools are incentivised to cream-skim as a result of institutional structures and the nature of demand for school places requires a detailed discussion of the nature of the utility function for the various agents in schools; this is set aside until the analysis of competition in Chapter 6. The purpose of this chapter is to examine the theoretical basis for arguing that parental school choice strategies differ by social background and the implications of this for developing models of school choice to predict levels of school stratification.

The observation that parents' school choice strategies differ by social class can be witnessed and measured across different types of institutional settings. For example, studies of the impact of the 1988 Education Reform Act showed that families already advantaged were more likely to gain places at desirable schools than disadvantaged families (Conway, 1997; Levačić and Hardman, 1998; Reay, 1998). Burgess and Briggs (2006) show that free school meals eligible pupils attend worse schools than pupils not eligible for free school meals who live in the same street. In the US, Hastings et al.

(2006) have shown that low income families are less likely to exercise choice in a public school lottery system. And of course, where choice is only possible through house moves in a strict neighbourhood schooling system, it has been shown that children from low income and less-educated backgrounds experience poorer than average school peer groups (Black, 1998; Bayer and McMillan, 2005).

However, these empirical observations contribute little to our understanding of *why* this relationship between school choice and social class exists. Without an understanding of the household decision-making processes that contribute to this differential behaviour, researchers cannot predict the magnitude of the change in sorting that is likely to result from a particular policy implementation, with the result that it is not possible to devise the optimal policies that are likely to be successful in lowering social segregation and increasing educational equity.

The first part of this chapter draws on the English sociological choice literature, developed within a post-structuralist tradition, to describe how the middle classes and working classes interact with the school choice process. This is a descriptively rich literature that illustrates the complexities involved in the way choice is offered, understood and actualised within families, and its claims are largely backed by small scale qualitative and larger scale surveys of parents. However, it is not able to make strong *predictive* claims about the effect of policy on stratification. This is because, while it can demonstrate how working class families are disadvantaged in the English system, it has no means of assessing whether they are better or worse off than under policy counterfactuals such as strict catchment area systems versus more radical reform programs.

By contrast, economists are particularly well-equipped to build theoretical models to predict the effect of policy changes. The next part of the chapter introduces general equilibrium models of household decision-making, first under a neighbourhood schooling model and then with various school choice programs. Compared to the rich post-structuralist accounts of the household behaviour, these models appear overly simplistic: they do not necessarily contradict post-structuralist models, they simply ignore most of the nuanced descriptions of behaviour. These models tend to suggest that (some) lower income families benefit from choice reforms (at least they do so more than wealthy families), but this is an almost trivial claim since the models tend to rely on the income constraint of families as the main variable for producing stratification in the neighbourhood schooling system.

The chapter goes on to evaluate the extent to which economic models that bear little resemblance to how parents actually make choices may lead to incorrect inferences. Box (1979) famously said: '*all models are wrong, but some are useful*'. The extent to which simplifications might pose a problem depends entirely on the research questions

demanded of the model. One general consequence of simplification is omitting key parameters so that the models, when calibrated with real data, might lead to biased estimates of the parameters in the utility function (Goldstein and Pauly, 1981). More seriously though, there must be a concern that certain modelling simplifications and unrealistic underlying assumptions will produce incorrect inferences as to the *directional* effect of changes in school admissions on stratification.

In its evaluation of the current general equilibrium of school choice literature, this chapter argues that it is of little use in assessing whether recent choice-based reforms in England are likely to be successful in altering the sorting of pupils into schools. Specifically, economists (such as Hoxby, 2003a; Burgess et al., 2005b) who make predictive claims that choice *must* lower school stratification because it removes constraints from the low income families who are most constrained by a neighbourhood school model are doing so only contingent on a narrow set of assumptions within very reductionist theoretical models drawn from a literature with little applicability to the English choice system.

The chapter concludes by describing how economists might proceed in terms of creating models of English-style school choice that could usefully predict how the system might respond to policy interventions. It emphasises how different the English school admissions system is to current economic models of choice: partly because parents are forced to purchase a home in the presence of considerable uncertainty about the outcome of a subsequent school admissions process, and partly because of the widespread prevalence of over-subscription criteria that are based on individual attributes such as religious adherence and child's ability.

3.1 Post-structuralist models of parental choice

The purpose of this section is to provide a rich description of how the social class of families shapes the manner by which they interact with the school choice process. The sociological literature on parental choice, as developed by Pierre Bourdieu, Stephen Ball, Diane Reay and others, is capable of doing this, using small-scale qualitative studies for the purpose of building and illustrating theories of choice and social class. Although certainly not individualist in its approach, it is a particularly useful literature for economists to read and incorporate into their models of decision-making. This is because, while it recognises that individuals are thoroughly embedded in social relations, it accepts that actual behaviour takes place at the level of individual actions. So, the research tends to focus on the '*micropractices of social reproduction, and on the situated enactment of class skills, resources, dispositions, attitudes and expectations*' (Ball, 2003a, page 3).

Social class, as an attribute of individuals that is distinct from income and education level, is not a description widely used in economics, but it is central to the post-structuralist literature summarised here. There is a degree of fuzziness in the categorisation of class. Ball describes his focus as being class as '*an identity and a lifestyle, and a set of perspectives on the social world and relationships in it*' (Ball, 2003a, page 6). However, in the empirical work described below, class positions are analogous with social relations in economic life or, more specifically, from employment relations. So, the classification of social class is consistent with Goldthorpe's theory and measurement of social class, i.e. working class means blue collar or routine service sector jobs and middle class means skilled non-manual work (Goldthorpe, 2000).

3.1.1 Middle class advantage in school choice processes

An important part of class analysis is the Bourdieu framework of analysis that holds that each new choice is confronted with assets or capital (economic, social and cultural) to be exchanged or invested (Bourdieu, 1986). Thus, middle class advantage in the school choice process comes via economic capital (advantage in terms of financial assets and income); social capital (the social networks relied upon to inform school choice decisions) and cultural capital (the attitudes and knowledge, defined by those in power, which makes the education system a comfortable familiar place in which they can succeed easily). The key claim of the sociological literature is that social reproduction – the maintenance of power and privilege between social classes from one generation to the next – lies at the heart of middle-class existence (Savage et al., 1992); and that:

> The market works as a class strategy by creating a mechanism which can be exploited by the middle class as a strategy of reproduction in their search for relative advantage, social advancement and mobility.
>
> (Ball, 1993, page 117)

The middle classes clearly have many advantages in non-choice systems of school allocation, and this post-structuralist literature does not attempt to measure the advantages of the middle class in choice systems relative to non-choice alternatives. Instead they provide a description of the two processes that lead to middle class advantage over working class families within a choice-based system. First is the suggestion that choice has different meanings in different class contexts because it is a socially and culturally constructed phenomena (Gewirtz et al., 1995). This means that families of different social class backgrounds engage in the choice process differently. Second, and building on the first argument, the choice mechanism that exists in England has been constructed in such a way that it requires time, effort, expense and skill; i.e.:

> [Resources and capital] that are unevenly distributed across the population but with which the middle class are particularly well endowed. The education market with all its risks is well accommodated to the dispositions and interests of the middle class.
>
> (Ball, 2003a, page 173)

Thus, this literature recognises that choice policies do not emerge from a social vacuum, and Ball (2003a) argues that these education policies are developed by the middle class to pursue their own advantage and interest.

A good example of this is how the middle class draw on their skills and social capital to gain information on school performance and admissions policies. This is not just through informal social networks. West and Pennell (1999) show that high socio-economic groups appear to have better information on, and understanding of school performance via league tables. Coldron et al. (2008) also report a social gradient in terms of the use of formal sources of information regarding schools. Mothers who had qualifications at level 4 or higher (degree level) were found to be three times as likely to use formal sources as those who had no qualifications. Also, families with parents who were in employment were more likely to use formal sources than either lone parent families or two-parent families where one or both parents did not work. They also report that internet access was important to a family's ability to access information about schools, which may place lower income families at a disadvantage in the process. Overall though, the most highly valued information was obtained informally, through school visits and talking to other parents and staff. Thus, parents with stronger social networks were more likely to gain more useful information about local school quality and admissions.

Coldron et al. also document the explicit means by which the cultural, social and economic capital of a family can be used to maximise chances of securing a place at a preferred school. They report that 8% of parents with children now at state maintained secondary schools admitted to coaching children for entrance tests; 5% reported ensuring their child was in the correct feeder primary school; 4% reported paying for extra tuition; 4% reported arranging extra-curricular activities; 3% reported moving or renting a house in the correct catchment area; 1% reported joining a church or place of worship; and 1% reported asking someone with influence in the process to recommend the child.

3.1.2 Middle class judgements of school quality

The middle class are not just characterised as being advantaged in the process of securing a place at a good school, they also have a very specific notion of what constitutes

'good', and this is crucial to understanding why markets are likely to become stratified. Ball characterises middle class choice as being dominated by maximising the child's likely exam results where '*the school is not represented as an independent variable with qualities of its own separate from its intake*' (Ball, 2003a, page 169). In other words, they have a narrow conception of quality as being analogous with expected test scores, and believe that the quality of school largely depends on the school peer group, almost irrespective of the quality of instruction.

The characteristics of a child's peer group at school have been shown to be an important input into an education production function, that also includes teacher quality, the quality of the headteacher and the financial resources or assets of the school (Robertson and Symons, 2003; Dills, 2005; Ammermueller and Pischke, 2006).[1] The interpretation of a 'peer effect' need not be restricted to the direct impact of a child and their family on other children's education achievement: we may also think that resources such as teacher or headteacher quality are sorted among schools according to school composition. Thus, a family seeking to maximise their child's achievement will, *ceteris paribus*, seek the school with the superior peer group. They may also be using the observed peer group of the school as a proxy for school quality given a complex education production function. This would be consistent with the observation that consumers in markets for complex goods and services over-rely upon simple accessible signals.

Survey evidence shows that parents report valuing qualities of schools beyond the basic peer characteristics of the school. However, from a modelling perspective, proxying parental perceptions of school quality using peer characteristics may be sufficient if the correlation between the two is reasonably high. Furthermore, it is important to interpret all of the reported surveys somewhat cautiously because survey responses may not give an accurate estimate of parents' true preferences. For example, stated preferences may be altered to fit social norms, emphasising a high value for education quality and child's happiness and potentially downplaying concerns for a school's social and racial composition.

There is considerable variation between the different studies in the responses of parents, as shown in Table 3.1, although proximity, discipline, exam results and child's preferences all regularly feature. The variation in responses is, perhaps, not surprising as the studies differed considerably in their methods. For example, some asked parents to choose reasons from lists while others used open interviews; the social class make-up and geographical location of the samples were different; they were carried out at different points in time, and so on. It is noteworthy that the most recent survey claims few parents (36% of all those who looked at performance tables) report an interest in value-added scores, arguably the most accurate published measure of school quality

[1] Ammermueller and Pischke (2006) estimate primary rather than secondary school peer effects, but the study is of interest because it distinguishes between school and classroom level effects.

(Coldron et al., 2008). The majority (80% of all those who looked at performance tables) are interested in the GCSE/A-level results of the school, even though the overwhelming determinant of these is not the quality of instruction but instead is prior attainment of intake and therefore the quality of peers.

3.1.3 The complexity and risk of choice for the middle class

Although the sociological studies claim the market favours the middle class, they also report a pessimistic view of the process of choice from this group. By constructing a belief that a child's educational success is crucial to their maintenance of social position, and in believing that educational success is largely determined by school peer group quality, the risks for the middle class associated with choosing an inappropriate school might be low probability but high consequence.

The risk in the choice process emerges from 3 sources:

- the risk from probability of not securing school of choice;
- uncertainty about actual school quality; and
- risk/uncertainty about the child's likely outcome within that setting.

Much of this risk arises because:

> Knowledge is increasingly indeterminate and contingent. Complete and completely reliable information may be an impossibility in human processing institutions, and the perceived adequacy or reliability of information available both undermines parents' trust of pay and expert actors and exacerbates anxieties.
>
> (Ball, 2003b, page 166)

He describes this as leading '*to a collecting, sifting, advertising and hearsay. In a sense, all information is considered but none, or almost none, is trusted*' (Ball, 2003b, page 167).

Thus, the economics perspective that individuals can calculate expected school quality based on a set of known probabilities is perhaps unrealistic, even as a representation of middle class decision-making. Crook argues that '*information overload leads to arbitrariness and necessary imcompleteness of even the most assiduous individual risk calculation*' (Crook, 1999, page 180). Because it is not clear what the 'right' choice is, for middle-class families school choice is '*a matter of uneasy compromises. Indeed, there is currently a kind of moral panic around schooling and school choice, particularly in metropolitan settings*' (Ball, 2003b, page 165).

Table 3.1: UK surveys of parents regarding choice of school

Author	Sample	Findings
Elliott (1982)	Ex-post interviews with parents at one secondary school	Proximity (1st); children given an important say (2nd); providing a balanced all-rounded education (3rd)
West et al. (1984)	Ex-post interviews with 216 Inner London parents	Proximity (33%); reputation (31%); relations went there (29%); good discipline (23%)
Alston (1985)	Ex-ante interviews with Inner London parents in primary schools	Child's view (65%); proximity (53%); facilities (48%); discipline (48%); exam results (38%)
Stillman and Maychell (1986)	Large scale ex-post survey across 4 LEAs	Academic record (1st); good discipline (2nd); proximity (3rd)
Adler et al. (1989)	Ex-ante survey with list of 4 items	Thought child would be happiest there (1st); child's preference (2nd); discipline (3rd); proximity (4th)
Coldron and Boulton (1991)	Ex-ante survey of 222 families asked for most important factor	Proximity (32%); sibling attends (16%); friends attending (15%); child's choice (10%); best educationally (10%)
Hunter (1991)	Ex-post interviews with Inner London parents	Good discipline (47%); proximity (42%); good exam results (39%)
West and Varlaam (1991)	Ex-ante interviews with Inner London parents	Child's choice (71%); good discipline (67%); good exam results (54%); proximity/access (53%)
West et al. (1993)	Ex-ante interviews with 70 Inner London parents	Exam results (33%); atmosphere/ethos (33%); proximity (25%)
Flatley et al. (2001)	Ex-post and ex-ante surveys of c. 3,000 parents	Academic outcomes (43%); siblings or proximity (40%); child's choice (31%); ethos (15%); quality of staff (14%); facilities (13%); behaviour (10%)
Bradley and Taylor (2007)	Ex-post (yr 9) interviews with c. 3% of secondary aged pupils	Views of family and friends (66%); school's location (63%); exam performance (38%); religion (8%)
Coldron et al. (2008)	Ex-post telephone survey of 2,215 parents in 59 schools.	Good reputation/Ofsted (40%); good exam results (33%); child's choice (31%); sibling goes there (28%); facilities are good (22%); friends are going there (20%); local school (20%); discipline/behaviour (19%)

3.1.4 Working class values and school choice

By contrast, according to Reay and Ball, '*working-class patterns of educational choice are characterised by ambivalence, and appear to be as much about the avoidance of anxiety, failure and rejection as they are about choosing a good school for my child*' (Reay and Ball, 1997, page 93). There are two parts to this argument. First, the ambivalence is consistent with working class families viewing their child's characteristics as 'fixed' and not susceptible to school effects. According to the school effectiveness literature there exists variation in school quality of the order of 8-15% of total variance in attainment (Reynolds, 1992, page 70), which means this perception is incorrect. However, the working class probably underestimate the extent to which school matters only as much as the middle class may overstate its importance in determining educational success.

The second part of the argument emphasises the contradictions and compromises in making choices because of the potential negative consequences of entering a middle class school.

> Such choices could set working-class children to fail in individualised, publicly humiliating ways in predominantly middle-class, high-achieving schools as opposed to the more masked, shared processes through which they fail (or are relatively successful) in local, inner city comprehensives.
>
> (Reay and Ball, 1997, page 97)

Rather than choosing popular, high reputation schools, working class preferences are strongly shaped by the parent's own (often negative) experience of school (Woods, 1993). They tend to value the accessibility and friendliness of teachers and rely on 'gut feeling'/intuition or favouring a sense of 'being at home'. Indeed, often working class parents were impressed when schools gave positive attention to less academically inclined pupils rather than focusing primarily on able students (Reay and Ball, 1997). Surveys show that in lower social class families the child's wishes are often decisive while for middle class parents the child's input into the process is limited (Coldron and Boulton, 1991; Ball, 1993). The result is that within middle class norms, working class families may appear to be 'bad choosers', but this is entirely a cultural judgement (Reay and Ball, 1997).

Several US studies are able to robustly show that low-income parents place lower values on academic characteristics when choosing schools (Fossey, 1994; Armor and Peiser, 1998; Schneider and Buckley, 2002; Jacob and Lefgren, 2007; Hastings et al., 2006). This is broadly consistent with the UK survey literature described earlier. However, Coldron et al. use their survey findings to temper the 'deficit model' of parents from lower socioeconomic groups being disadvantaged in the choice process, as is preva-

lent in this literature that contrasts between 'skilled choosers' versus 'disconnected choosers' (Gewirtz et al., 1995) and 'alert' versus 'inert' (Echols and Willms, 1992), stating:

> It is a widespread assumption that parents from lower socioeconomic groups are being denied access because they are less able to understand the admissions process and therefore less able to successfully negotiate it. We found no evidence to support this. While more educated parents were likely to access more information very few parents felt they were lacking basic information about secondary schools and there was no evidence that parents who were less educated had any reduced chance of gaining their first preference.
>
> (Coldron et al., 2008, page vii)

However, it is not clear that this should be taken as contradicting the work of Diane Reay and others: Coldron et al. simply confirm that the working classes do not want to engage in the school choice process in the same manner as the middle classes do, and that they have different aspirations in terms of schools they wish their children to attend.

3.2 Economic models of parental choice

In contrast to the sociological literature presented, the economic paradigm is limited in its ability to explain why strategies differ by social class because it lacks a proper concept of 'class' and in particular is incapable of explaining how and why social class norms and practices emerge. However, it can incorporate systematic differences in human behaviour into decision-making models in a variety of ways. In the economic models currently prevalent in the school literature, all parents act in a rational manner, making a choice of school from a well-defined choice set with the goal of maximising household utility based on fixed preferences and subject to a budget constraint. Within this framework, it is possible to model parents as differing according to social class, or socio-economic status, which is usually classified as being analogous to differences in human capital, measured using parental occupations or (more usually here) parental education levels.

The economic models of parental decision-making, described in some detail for the remainder of this chapter, appear very reductionist and overly simplified compared to the rich descriptions of human motivations and behaviour used by post-structuralist sociologists. There is nothing intrinsically wrong with this process of simplifying descriptions of society within economic models; indeed it is generally necessary in order

to build models capable of predicting the effects of a change in policy. However, the risk is that simplifying assumptions substantively alter the predictions the model makes, and so the reduced model causes us to make incorrect inferences as to the effect of a policy implementation were it to take place in the real world.

One criticism commonly made of these models is that they contradict the wide body of experimental evidence demonstrating that human behaviour deviates in systematic ways from the idealised behaviour attributed to expected utility maximisers (Tversky, 1996). However, this observation is not enough to dismiss economic theory building. Roth (1996) argues that most economists view the rational model as a useful approximation, rather than as a precise description of human behaviour. In fact, arguably the school choice decision is better described in terms of this economic decision making than so many other decisions in life: it is usually a conscious, deliberate and considered decision, made in the presence of considerable information. But it is worth noting that there is a growing attempt by economists to move away from an overdependence on idealised models of hyper-rationality. The emergence of agent-based computation (described later in the chapter), which is able to explicitly model myopic behaviour, satisficing, and so on, within the field of economics is one example of the development of new tools to facilitate this.

There are three distinct explanations economists are able to give for why parents of different social classes choose different schools, described below.

3.2.1 Income constraint

In modelling social class differences in school choice decisions, the first perspective is that the tighter budget constraint of low income families is sufficient to explain their inability to purchase houses next to popular schools, pay for private tuition for entrance tests, or take long journeys to school. Under this account, the underlying preferences of lower and middle class families for schools are no different. The policy implications of this would be that social stratification between schools could be significantly lowered by reducing the costs of accessing schools for low income families, for example, by designating places at schools for out-of-catchment children or by subsidising transport to school. This type of explanation for class differences is similar to that of sociologists Breen and Goldthorpe (1996), who suggest that class differences in educational experiences and outcomes are a reflection of the objective differences in the material situations of families in different class situations.

However, the econometric estimates of house location decision suggest this explanation is not sufficient on its own to explained observed school stratification. Bayer and McMillan's work takes an area (San Francisco) where choice of school is only possible

by moving house because strict residence requirements are in place and estimates a random coefficient, discrete choice model of elementary school 'choices' by parents. They find significant evidence for differences in preferences in the sense that more educated parents are willing, *ceteris paribus*, to spend more for high quality schools (Bayer, 2000; Bayer and McMillan, 2005). This finding is confirmed by Nesheim (2002), who uses a similar approach to estimate a correlation between parental education and willingness to pay for school quality (given income) of 0.59.

3.2.2 Shape of the utility function

The second explanation given by economic models is that preferences for school quality differ by social class, with the utility derived from greater school quality being lower for low social background families. This might manifest itself in many ways. The family may be less willing to substitute consumption for school quality 'purchased' through the housing market or transport costs. Alternatively, the family may calculate that the utility gain from superior school quality is not enough to offset utility loss from longer journeys to and from school. Furthermore, other aspects of the school environment may enter the utility function, and they might place significant utility on the child's own expected happiness at secondary school, thus favouring allowing their child to continue to be educated with friends.

The post-structuralist theories described in the previous section argue strongly that utility functions do differ by social class in this manner. However, the problem with these types of 'explanations' within an economic model is that they fail to explain where these preferences come from. This is because economists treat social norms and culture as arbitrary preferences of the individual and thus provide no mechanism for explaining how they emerge and change over time. This limitation may be crucial in the context of education policy. For example, it is possible that the publication of league tables in 1992 altered some families' perceptions of school quality, leading them to place a great emphasis on GCSE results than (for example) behaviour or ethos in a school. Thus, a policy intended simply to improve the quality of information in the marketplace may have altered the nature of some families' utility functions. As a further example, given that we know social networks are crucial in the transmission of information on school quality, it is also possible that conversations with social neighbours can alter one's own willingness-to-pay for school quality or other aspects of the utility function. Thus, a family's utility function is partially shaped by their neighbour's utility function.

3.2.3 Informational advantages

The third explanation for differences in school choice strategies is that higher social background families are endowed with an informational advantage (from social networks, for example). This manifests itself in several ways. First, it gives them better measures of school quality. Second, they select a school from a larger choice set. Third, they have a superior ability to estimate the probabilities of acceptance at different schools, which is particularly important under a First Preference First choice system. The empirical evidence in the previous section suggests all these advantages are likely to be present. However, this explanation is, again, necessarily partial because we need a model of the costs of obtaining information for different groups, relative to their own perceived benefits in holding this information. It is perfectly possible that working class families place less value on this information, thus information differences actually result from differences in the shape of household utility functions.

3.3 General equilibrium models of location choice

Having shown the three ways by which economists can incorporate social class differences in household decision making, the remainder of the chapter now examines actual economic models of school choice in detail. The purpose of this section is to provide an overview of existing general equilibrium models of location choice where school assignment is decided strictly via a residence requirement restriction and thus households purchase homes and public school access as bundles. This is the traditional means by which parents 'choose' schools in the United States and is similar in nature to catchment areas or proximity oversubscription criterion in the UK context. The models discussed are all hedonic pricing models, following those first developed by Tinbergen (1959) and Sattinger (1980). A central feature of the models is that peer groups, and thus school quality and the value of housing, are endogenously determined. The models match consumers to locations and find prices that separate people based on willingness to pay for locational quality, of which local school quality is one dimension. Econometric studies have consistently shown that parents are willing to pay for school quality through the housing market (Black, 1999; Bogart and Cromwell, 1997, 2000; Goodman and Thibodeau, 1998; Sieg et al., 1999; Gibbons and Machin, 2003; Cheshire and Sheppard, 1995; Leech and Campos, 2003; Rosenthal, 2003).

The hedonic pricing models are, as with all models, necessarily reductionist. They simplify the consumer decision-making process and utility function; they emphasise homogeneity rather than heterogeneity of households and location characteristics; and they analyse outcomes solely in equilibrium. The advantage of these models are the strong predictions they are able to make about the effect of policy changes. However, they

treat global regularities (concepts that are not reducible to micro-phenomena) such as social norms and culture as arbitrary preferences of the individual. The methodological individualist approach where all social phenomena are explained in terms of individual actions, need not understate the importance of these global regularities; it simply provides no mechanism for explaining how they emerge and change over time.

All models of location choice with neighbourhood schooling are some variant on the Epple and Romano (2003) model. They set up a simple locational choice model for households, with entry to schools restricted by a neighbourhood residence requirement and all students attending a public school. Overall land capacity in Epple and Romano's economy is constrained to the level of the population (relaxing this assumption introduces additional equilibria) and neighbourhoods are equally sized to avoid multiple equilibria. This seems a reasonable starting assumption in the UK context, where land in towns is highly constrained (a more complex approach would be to follow Epple and Platt (1998) who have a house supply function that allows a community to grow). Papers containing developments on the Epple and Romano model that are referred to in the remainder of this chapter are summarised in Table 3.2.

Table 3.2: Summary of general equilibrium of household location articles

Author (Date)	Key contributions	Individual characteristics	Utility function contains	Model estimated?
Epple and Romano (2003)	Comparisons of neighbourhood schooling models to open enrolment models. The model allows households to vote on local public spending levels.	Income; child's ability	Consumption, educational outcome (a function of per pupil spending, mean ability, child's own ability).	No.
Nechyba (2003a)	Emphasises role of innate house and district characteristics in explaining sorting levels.	Income; child's ability	Consumption, school quality (a function of per pupil spending and mean ability), housing type and exogenous district characteristics.	No.
Nesheim (2002)	Theoretical models to show the conditions under which an education production function can be estimated without bias.	Parent's education; income; child's ability; taste for schooling.	Consumption and utility from expected educational outcome (a function of taste for schooling, mean parental education in district, own parental education, childs ability).	No
Ferreyra (2005)	Use of a full-solutions method to estimate her model. Empirical implications of two different types of voucher programs.	Income, religious type	School quality (a function of mean parental income and spending per student), consumption, housing quality (size, age and amenities), idiosyncratic preferences for location and school type.	Yes, with full-solution method minimising a well-defined distance between the predicted equilibrium and the observed data.
Bayer et al. (2004)	Two-stage estimation of model. Shows large effects of school quality changes on changes in stratification.	Parent's education; income; race; employment status; household composition.	Housing type (age, size, type, tenure); neighbourhood amenities (school; crime level; land use; topography); neighbourhood mean parental education; house price; travel-to-work distance.	Yes, via a 2-stage solution holding observed price and location decisions of others constant and boundary fixed effects in 2nd stage.

3.3.1 Modelling the household utility function

The elements and shape of the household utility function are the key determinants of the nature of stratification and house price differentials in equilibrium. In the basic Epple and Romano model, preferences of households are assumed to be fixed, meaning that there is no mechanism by which preferences can change in response to their environment, and households have only two attributes: household income, y, and child's ability, b, with a correlation between y and b that is either positive or zero. Households purchase a dwelling with annualised amortised house price, p, and derive utility from consumption (income after tax and housing costs) and final educational achievement of the child, a. Education achievement is the outcome of an education production function containing initial ability (which we can take to mean measured test score at age 11), b, and school quality, q. School quality is in turn a function of the per pupil expenditure, X, and the average ability of peers in the school, θ.

Epple and Romano model the peer effect in district j as mean ability of peers, in-line with much of the existing empirical literature (Feinstein and Symons, 1999; Ammermueller and Pischke, 2006). However, there are alternative models where it is a function of parental income or education (e.g. Fuchs, 2005; Nesheim, 2002); others claim peer effects are non-linear (e.g. Summers and Wolfe, 1977). It matters how we choose to model (parents' perceptions of) peer groups in terms of equilibrium outcomes. Within a simple model of neighbourhood schooling that predicts complete income stratification, the resulting variance in school quality would be highest where an income peer effect is modelled and lowest with the child's ability peer effect.

In the UK context it is unnecessary to model Tiebout sorting whereby a one-school district can vote for a tax and raise revenue directly to fund that school's expenditure. Voters can, and do, also influence the level of education funding in the local authority, but the magnitude of this funding difference is very small compared to the variation in per pupil funding distributed to schools to compensate for educational and social disadvantage (Levačić et al., 2005). This differential per pupil funding and its effect on school quality can be treated as dampening the effect of the peer group, θ, which is redefined as entering the educational production function as the totality of the direct effect of the peer group on school quality less the indirect effect of the differential per-pupil funding.

The household utility model (often with Cobb-Douglas form imposed) can be generally written as:

$$U(b, y) = U(y(1 - t) - p, a(\theta, b)) \qquad U_1, U_2, a_1, a_2 > 0$$

3.3.2 Equilibrium outcomes

General equilibrium models make end-state predictions that are computed or solved mathematically and are independent of starting circumstances of the model. These models always have at least one equilibrium because there always exists the trivial (unstable) equilibrium where all districts have identical peer groups such that there is no gain from moving district (Nesheim, 2002). Nesheim shows that for his utility function that is linear in consumption, there is a separating equilibrium (i.e. an equilibrium where peer groups are not identical) if and only if the correlation between the attribute that produces the peer group (in his case parental education) and willingness to pay for school quality is greater than zero.

In the Epple and Romano equilibrium there exists a strict hierarchy of house prices and the allocation exhibits boundary indifference and strict household preference within boundaries for their own neighbourhood. The elasticity of the price premium with respect to neighbourhood quality is solely a function of:

1. the elasticity of child's schooling attainment with respect to school quality (where school quality is solely determined by mean child's ability in the school);
2. the correlation between willingness to pay for school quality and child's initial ability, b.

If the first elasticity is large, then large proportionate differences in price are needed to segregate people. If the second correlation is high then a small difference in price is needed to segregate people.

The key outcomes of interest for these chapters are always the levels of income and ability stratification in the model and, as a consequence of these, the extent to which quality varies across schools. Epple and Romano use a somewhat non-standard definition of stratification that is renamed to maintain consistency with later chapters:

Partial income stratification: for any given ability child, only higher income household pupils attend higher quality schools.

Partial ability stratification: for any child with given household income, only high ability pupils attend higher quality schools.

The equilibrium properties of the basic neighbourhood schooling model are shown in Table 3.3. The first row shows the equilibrium outcomes where child's ability is assumed to be uncorrelated with household income. In this situation, schools would not become stratified, provided that willingness-to-pay for school quality was not related to the child's initial ability. The second row shows the equilibrium outcomes where income and ability are positively correlated. In the absence of Tiebout sorting, variation in school quality, q, is entirely driven by differences in the mean ability of

peer groups between schools. Even without partial ability stratification, differential school quality can result, but only if household income and child's ability are positively correlated.

Table 3.3: Equilibrium outcomes of neighbourhood schooling model

	$\delta(\frac{\delta U/\delta q}{\delta U/\delta y})/\delta b = 0$	$\delta(\frac{\delta U/\delta q}{\delta U/\delta y})/\delta b > 0$
$E[b\|y]$ constant	No partial income stratification No partial ability stratification q constant across schools	Partial income stratification Partial ability stratification q varies
$E[b\|y]$ increasing in y	Partial income stratification No partial ability stratification q varies	Partial income stratification Partial ability stratification q varies

It is impossible to describe the qualitative nature of the equilibrium prices and sorting in more detail without imposing some assumptions about the distribution of parameters and the functional form of the utility function on the model. For example, the distribution of income and child's ability are crucial determinants of the shape of the equilibrium because they determine both the relative supply of consumer types and relative demand for locations in the economy (Nesheim, 2002).

3.4 Models that predict incomplete stratification

The simplicity of the basic Epple and Romano model allows the relationship between the model parameters and the equilibrium properties to be clearly understood. However, the basic model described predicts complete income stratification between schools under a neighbourhood schooling system, meaning (for example) that rich families live in one district entirely separated from poorer families. This is clearly not consistent with real world data from the US or the UK. For example, in areas of the UK where most pupils attend their catchment area school and there are no own admissions schools, free school meal segregation is only around 0.2 on the index of dissimilarity.

This section introduces potential modifications to the model that are intended to allow it to mimic the real world social processes that generate incomplete income stratification across neighbourhoods. First, a set of new parameters are introduced that make this possible within a general equilibrium framework. While many different approaches are discussed, they are not mutually exclusive, and it is likely that the real world displays all these features to some extent.

Following these new general equilibrium parameters, a set of further explanations for

incomplete income stratification across neighbourhoods are given, all of which are either very difficult or impossible to incorporate into a general equilibrium framework. To overcome the modelling restrictions of general equilibrium, agent-based computations models are introduced as a more flexible and dynamic approach to modelling the house purchase decision.

3.4.1 Within a general equilibrium framework

Three new types of parameters are discussed here that have the potential to generate incomplete income stratification in a neighbourhood schooling model. The plausibility and usefulness of each new parameter can be evaluated in two ways. First, does the existence of the parameter correlate with a real world phenomena that we know to be present in the UK context through evidence presented earlier? Second, does the parameter generate the patterns of price differentials and income stratification that we know exist in parts of the UK where catchment areas dominate?

Heterogeneous housing and location characteristics

Epple and Romano assume that housing plots are perfectly malleable in the long term; so all plots within a neighbourhood command an equal price, therefore differential school quality is the only parameter driving between-neighbourhood price differentials. However, at least in a UK context, it seems important to treat housing plots as heterogeneous, even in the very long term (given planning restrictions). They are uneven in size, so offering different square footage, number of bedrooms and sizes of gardens, and so on. Nechyba's set of general equilibrium models of location choice incorporate 15 different housing types, h, that can be placed in the utility function as one dimension of location characteristics, l (Nechyba, 2000, 2003a,b, 2004).

Location characteristics will also differ, even within a school catchment area if families derive utility from particular locations for reasons that are unrelated to schooling. If an entire school district or catchment area is desirable, for example, because of closeness to a railway station, motorway or amenities, this would only serve to enforce complete income stratification and lead to a unique equilibrium in the neighbourhood schooling model. However, if parts of a neighbourhood have desirable amenities but the rest does not, this would produce heterogeneous within-neighbourhood prices. These location amenities must be those that can be treated as fixed and therefore exogenous to the assignment problem. Railway stations and motorways would be good examples of this; restaurants and shops would not since their own location decision depends on local neighbourhood family attributes.

Location amenities, r, and housing types h would then both enter the utility function as follows:

$$U(b, y) = U(y(1 - t) - p, a(\theta, b), l(h, r))$$

However, this specification would imply that all individuals placed the same utility on each amenity. In reality, we probably believe there exists systematic variation in willingness to pay for amenities across individuals. Furthermore, it is unlikely that all individuals rank the desirability of amenities identically (for example, some families would prefer to be closer to the motorway than the railway station, and vice versa). Indeed, introducing both these adaptions to the model would again lower income stratification.

Characteristics of neighbours

If households derive utility from a neighbourhood effect of families living in close proximity, say within a sub-district, we could allow this extra parameter to enter the utility function directly. Neighbours might be valued for social network reasons, or because of the local amenities they attract. We could assign each school district, j, k sub-districts, which would need to be non-overlapping to make the equilibrium solution tractable. Neighbourhood effects, n_k, might be represented as the mean income, y, of households in the sub-district, or some other function, and would enter the utility function as follows:

$$U(b, y) = U(y(1 - t) - p, a(\theta, b), l(h, r, n_k))$$

However, while this specification could explain within-district price heterogeneity on its own, it does not explain incomplete income stratification as required. Nevertheless, it is still important to form a view of the extent to which the neighbourhood peer group enters the utility function because this parameter affects the extent to which school choice models reduce residential stratification.

Variation in demand for schooling

The basic neighbourhood model described in the previous section does allow demand for schooling to vary across individuals. First, the willingness to pay for school quality via house prices is always modelled as increasing in income y. Second, in some specifications the willingness to pay for school quality increases with child's ability b. The first assumption – that schooling is a normal good – seems relatively uncontroversial. However, it is less clear that it is desirable to model the second. This implies that parents of higher ability children value a good peer group more than those parents with

lower ability children. This would be the case, for example, if the capacity to benefit from a more able peer group is greater for higher ability children. The only econometric study to support this modelling approach is Hastings et al. (2006), who use a natural experiment created by the re-zoning of school catchments following the dismantlement of bussing to show that preferences for measures of academic achievement are increasing in both income and baseline academic ability.

Given limited evidence for preferences varying by ability, the following two approaches to modelling variation in demand would seem to be more relevant. If we believe that demand for schooling varies by household 'taste', τ, (where τ and y are not perfectly correlated) we can achieve an incomplete sorting of households by incomes. Epple and Platt (1998) model this for a general (non-school) location decision. Nesheim (2002) also incorporates a preference for schooling parameter, and describes it as the level of parental altruism or parental perceptions about their child's expected return to education and the effect that differential school quality is likely to have on a child's education outcomes. Once again, adapting the basic model:

$$U(b, y) = U(y(1-t) - p, a(\theta, b, \tau), l(h, r, n_k))$$

If $\tau_1 < \tau_2$, we can interpret this as meaning parent 1 underestimates the effect of school quality on their child, relative to parent 2; *or* that parent 1 correctly identifies that their child would not benefit as much as parent 2's child from higher quality education; *or* that parent 1 places less value on the final educational achievement of their child.

An alternative modelling approach to variation in demand for schooling is to assume that some households value certain aspects of schools that are unrelated to pupil outcomes. For example, some parents might particularly care about the child's enjoyment of school; strict discipline; sports or music facilities. Alternatively, it might be the case that all households solely maximise child's outcomes, but have different beliefs about the best means to achieve it. For example, Maria Ferreyra (2005) models idiosyncratic preferences for religious schools. We can model this diversity of preferences for school types (assuming they have no impact on school quality) as a characteristic of the school in district j, with households having idiosyncratic preferences, s_j, for these characteristics:

$$U(b, y) = U(y(1-t) - p, a(\theta, b, \tau), l(h, r, n_k, s_j))$$

3.4.2 Limits to general equilibrium modelling

The general equilibrium models can achieve incomplete income stratification in mathematically tractable models, as we require, but we may not believe the assumptions they make reflect the true causes of this phenomenon. This section explores a set of issues that might lead us to favour a computational, rather than an analytic, solution to our model. A computational model is still simply a formal representation of a theory about some empirical system, but it is usually a set of algorithms in a computer program with model behaviour explored by running a program over a parameter space.

Some issues described below can be dealt with within a computational general equilibrium (CGE) approach. CGE takes a set of equations and computes equilibrium solutions for a set of parameter values. It is therefore a similar, but more flexible, approach to traditional deductive general equilibrium modelling. Maria Ferreyra has built CGE models of school and neighbourhood choice. Those modelling assumptions that cannot be addressed within a CGE approach can be incorporated into the more flexible agent-based computational models described in the next section.

Complex neighbourhood effects

In the previous section utility models were described where households derived a utility from living with neighbours in a sub-district, where sub-districts are not overlapping and are nested within districts. CGE models can deal with this type of sub-district, as shown in Ferreyra (2005). However, suppose we wanted to make the concept of 'neighbours' more complex so that each household had a uniquely defined set of neighbours. This set might be the characteristics of the immediate neighbours in houses contiguous to theirs (known as the 'Moore' neighbours) or a complex decay function of decreasing utility placed on the characteristics of households the further away they live.

The consequences of this 'overlapping' perspective on neighbourhoods are two-fold. First, each cell or house has to be modelled as having unique characteristics, so households choose houses and not school districts. This requires a modification of CGE algorithms such as those used in Ferreyra (2005). More fundamentally though, these types of neighbour preferences ensure interdependent decision making that is capable of setting up an out-of-equilibrium dynamic whereby sorting emerges from an initial random allocation. Schelling (1971, 1978) shows how stark patterns of segregation emerge from local migratory movements among two culturally distinct, but relatively tolerant, types of households. In fact, even where households have strict preferences for integration, segregation still emerges (Pancs and Vriend, 2003). Thus, a Pareto optimal equilibrium can never be achieved due to its inherent instability (so CGE is unworkable with these types of preferences).

There is empirical literature that supports the needs for more complex modelling of neighbourhoods. The literature on the UK housing market stresses that the behaviour of house prices appears to be complex and non-linear as a result of preferences for neighbours. Galster (2002) and Galster et al. (2000) stress the importance of thresholds as one form of non-linear behaviour: neighbourhoods do not start to decline or gentrify until they pass some trigger-point. But, once neighbourhoods pass a certain threshold, their character can very quickly change (Meen and Meen, 2003).

Bounded rationality and psychic costs of moving

Some might argue that, rather than households choosing houses on the basis of hyper-rational decision-making (optimising a utility function with precise information and cognitive ability), households actually 'satisfice' when making choices. Satisficing refers to a behaviour which attempts to achieve at least some minimum level of a particular variable, but which does not necessarily maximize its value. The word satisfice was coined by Herbert Simon (1947) who pointed out that human beings lack the cognitive resources to maximize: we usually do not know the relevant probabilities of outcomes, we can rarely evaluate all outcomes with sufficient precision, and our memories are weak and unreliable. A more realistic approach to rationality, known as bounded rationality, takes into account these limitations. Satisficing can actually be modelled as an optimisation problem for a household, where all costs, including the cost of the optimisation calculations and the cost of getting information for use in those calculations, are considered.

There are, of course, very real financial costs associated with the transaction of a house purchase (stamp duty, lawyers' fees and removal costs). However, another important cost (that it is quite difficult to monetise) is the psychic cost associated with moving away from the social networks developed in an area over time. A third cost is the information cost associated with gathering the information necessary to optimise. This includes school quality, neighbourhood amenity qualities and house prices for all houses in the economy. 'Satisficing' could therefore be a strategy that households use to balance search costs against potential gains and might explain why households would limit their search field to, for example, just the nearest 10 school districts. Meen and Meen (2002) report that 60 per cent of moves in England take place over less than 5 miles. However, the increase in information on the internet is substantially lowering costs of searching for both schools and homes. Coldron et al. (2008) report that two-thirds of parents find information on schools to be readily accessible, with internet sources used by many (including 44% accessing achievement and attainment tables, 29% reading Ofsted inspection reports and 25% looking at school websites). This represents a substantial change from the start of the 1990s, when a parent would often

have to write to each school individually to collect information on exam results.

The second type of cost described above – the psychic cost – can actually be modelled in a CGE model as an idiosyncratic preference for a particular district. However, the transaction costs associated with an actual move are fundamentally in conflict with the concept of an equilibrium analysis of house location decisions because equilibrium models are not dynamic models of house moving decisions and do not rely on initial locations. By contrast, a more realistic way to model the house purchase decision might be as a two-stage decision-making approach by households. In the first a (myopic) couple without school aged children choose where to live on the basis of neighbour characteristics, locational amenities, house quality, but ignoring school quality. In the second stage, the households need to send their children to secondary school. They then incorporate school quality into their optimisation problem and look at utility in all possible locations (in the economy or a closer area). They then decide whether to move or not, but only doing so if the benefits to moving outweigh the psychic and financial costs.

Disequilibrium and dynamic models

GE and CGE models are capable of telling us about the potential state of the world in equilibrium, but not at any other point. Equilibrium refers to a set of prices where the market 'clears', such that the houses supplied in each district match demand for homes, and specifically this literature uses a variation of a Nash equilibrium which takes place when no household can improve their utility, given other household's choices (Nesheim, 2002). Assuming an equilibrium state exists, it would only be of particular interest, over and above disequilibrium states, if the following properties held. First, the equilibrium is more interesting where is it unique. Uniqueness means that there is a single global equilibrium, with an absence of other local equilibria. Proving uniqueness in GE theory is never straightforward (e.g. Kamiya, 1984). However, in general, the greater the weight attached to the actions of others compared with one's own actions (i.e. the greater the importance of externalities), the more likely is the possibility of multiple equilibria. Second, the equilibrium should be stable. Where there are threshold effects, even though there may be a unique equilibrium with a given set of parameters, a small stochastic change in these parameters might 'tip' the model towards a completely different state. Third, the equilibrium should be achievable with rapid speed of convergence towards it.

This third condition is certainly not met in the context of housing and schooling markets. At any point in time, it is possible that the separating equilibrium is not actually obtainable giving the starting conditions. More likely, it is obtainable in theory, but the system remains out of equilibrium for long periods of time. Indeed,

where conditions in the market constantly change as a result of demographic changes, it is possible that equilibria are never reached. This is the situation in the US housing market, where Krupka (2008) has argued the data show that observed income mixing in cities is indeed unstable, but that the process of adjustment to a separating equilibrium is very slow.

Even if market conditions are held constant, we can only estimate the length of time until equilibrium is reached using a dynamic model of the actual house purchase process. By contrast, GE models employ the use of a Walrasian Auctioneer (or equivalent clearinghouse construct) that determines prices to ensure each market clears (Katzner, 1989). An actual housing market bears no resemblance to this since it has finite subsets of households willing to buy and sell at any particular point in time; the need for most households to buy and sell simultaneously, resulting in chains of house transactions; and local estate agents acting as negotiators of the local market, but not as clearing houses since they never hold property themselves. All this explains why global equilibrium may not be met in practice, or may take years or decades to emerge.

3.4.3 Agent-based computational models of school choice

Agent-based computational models are introduced briefly here as a particularly flexible and dynamic modelling strategy that can overcome the limitations of GE and CGE solutions described in the previous section. It provides a means of modelling repeated local interactions between a very large number of micro-agents – in this case households buying and selling houses. The term "agent" refers broadly to an individual or physical entity that is represented as bundled data and behavioural methods within a computationally constructed world and can be an individual, institution or physical entity (Tesfatsion, 2006). As opposed to traditional models that assume either a small number of dissimilar or numerous identical actors, agent-based models normally include large numbers of heterogeneous agents. Interactions between them give rise to global regularities such as peer groups, house price differentials and stratification. These global regularities in turn feed back into the determination of local interactions.

Agent-based systems are said to be complex, meaning that they are composed of interacting agents with system-wide properties (such as segregation) exhibited that arise from the interactions of the units and that are not properties of the individual units themselves (Flake, 1998). Agent-based systems are also said to be adaptive, meaning they include goal-directed agents that are capable of reacting to their environment and directing this reaction towards the achievement of built-in (or evolved) goals (Tesfatsion, 2006).

Agent-based computational modelling is a particularly flexible form of simulation. It can facilitate the modelling of complex interactive processes involving kinks, jumps, and other forms of discreteness imposed or induced by empirical constraints. It can also deal with bounded rationality, asymmetric information, strategic interaction, expectation formation on the basis of limited information, mutual learning, social norms, transaction costs, externalities, market power, predation, collusion and the possibility of coordination failure.

In agent-based simulations, the finite number of agents (households) are created with attributes (such as income, initial ability, tastes, as in GE models) and behavioural methods, which are rules for how they interact. For example, agents would need to be able to:

1. calculate their utility;
2. decide when and where to look for a new house;
3. negotiate a house price with another agent;
4. and move house.

Agent-based computational modelling requires the construction of dynamically complete economic models. Simulations are run from a set of initial conditions of house assignments and prices and must permit and fully support the playing out of agent interactions over time without further intervention from the modeller. Rather than try to calculate the equilibria states of the system, the focus is often on dynamics and transient trajectories. This all means that modelling can proceed even if equilibria are computationally intractable or non-existent. The stationery-state form tends to be emphasised rather than the equilibrium. In our case, this could be defined as the point of continual market clearing, of unchanging structure or of unchanging global regularities.

The literature on agent-based computational modelling in economics stretches back to Schelling's famous model where segregation emerges from residential location models with agents who have preferences for integration (Schelling, 1971). However, applications to school choice decisions are all very recent (and US-based) and are mentioned in the next section.

3.5 Decision-making under school choice programs

All models presented thus far have sought to make stratification predictions where school assignment is decided by strict residence requirement restrictions. Even under this very simple assignment criteria, models that adequately reflect the salient features

of households, locations and outcomes (i.e. heterogeneity of preferences and housing types both contributing to incomplete stratification in equilibrium) quickly become quite complex. This section looks at the implications of modelling greater complexity to the supply of schools via school choice programs. These models are important because in the absence of large scale real life experiments that alter school assignment rules, the models offer a clear framework for predicting the effects of school choice on student sorting.

Models of quite simple public and private sector choice programs that exist in the general equilibrium literature are described. Next, two attempts are made to model some aspect of choice using the agent-based computational approach described in the previous section. In both these fields, the literature is relatively recent and therefore under-developed, and in any case is entirely dominated by US researchers. Nevertheless, each model is assessed in terms of applicability to the English institutional structure.

3.5.1 Introducing school choice to general equilibrium models

There are two main developments in terms of modelling school choice within a general equilibrium framework. The first set of papers adopt models of choice in the public sector to a scenario whereby public schools make places available for non-neighbourhood families and these are allocated by lottery. The second set start with a strict neighbourhood schooling model, but allow households to use vouchers to purchase schooling in the private sector. The latter models are of limited applicability to the UK where private school vouchers have not recently been used (although Direct Grant school places and the Assisted Places Scheme were examples of this).

Universal lotteries for school places

Epple and Romano (2003) develop a basic model of school choice within a public school system. Schools have no neighbourhood residency requirement, no priority places for local residents and no capacity constraints. Given this, the cost of transport between districts is the key determinant of stratification. Where transportation costs are set to zero, all neighbourhoods and schools are perfectly integrated with identical house prices because schooling is entirely costless. Subsequent models introduce a simple flat rate inter-neighbourhood transport cost to a school in any other neighbourhood. Epple and Romano show that this once again creates stratified schooling (of the same nature shown in earlier sections where the residence requirement drives stratification).

For a separating equilibrium in a two neighbourhood example, house prices would be higher in neighbourhood 2 than neighbourhood 1 to reflect the higher school quality.

Some, but not all, households in neighbourhood 1 will choose to exercise choice and travel to school in neighbourhood 2. This model makes it clear that it is only the higher income families from the poorer neighbourhood 1 that will exercise choice where there are transport costs. The winners and losers from this type of choice system (relative to strict neighbourhood schooling) are complex. Higher income households find it cheaper to buy a house in neighbourhood 2 (but their house price would fall if they already live there). However, they would see some deterioration in their child's peer group. The lower income families in neighbourhood 1 unambiguously lose since they will see a deterioration in their local school's peer group, and it is not utility-maximising for them to pay the transport costs to neighbourhood 2. It is the group just below the median income that gain since they are now able to live in neighbourhood 1, yet access a higher quality school by paying the transport costs that are lower than the house price differentials between neighbourhoods 1 and 2.

The transport costs in these models can reflect both the disutility experienced by the child through long school journeys and the financial cost incurred by the household. Bearing this in mind, a flat rate transport cost is unrealistic in several respects. The cost of transport between any two districts depends on the distance between those two districts and public transport availability. This would have implications for patterns of stratification across a city. Also, we might want to argue that not all families face the same cost of transport. It would be higher where a family do not have a car available at that time of day, or where parents are working and cannot transport their children to and from school. If the cost of transport for any particular journey was decreasing in income, this would again serve to reinforce stratification.

Brunner and Imazeki (2005) add to the findings of Epple and Romano models by calculating who benefits from universal choice programs. They emphasise that the sizes of gains and losses under a universal voucher depends on changes in peer group composition and changes in housing values. In markets where Tiebout sorting was very high under the neighbourhood schooling model – meaning households actively selected homes located close to good schools – potential changes in peer group composition and therefore loss in housing values create an incentive for high income households to vote against choice reforms and low income households to vote in favour of reforms.

Vouchers for private and religious schools

The next set of choice models introduce a voucher program for private schools into a neighbourhood school system. Once again, it is Epple and Romano that produce the first model on which all others are based.

Epple and Romano (1998) contrast a neighbourhood schooling system of public schools with a model of competitive, tuition-financed private schools (with no public schools). In this model, private schools are able to charge differentially priced tuition based on the peer characteristics of students. The result is a strict hierarchy of school qualities with stratification by ability and income. In the private school model, high-ability, low-income students receive tuition discounts, while low-ability, high-income students pay tuition premia. The beneficiaries of this system are high-ability students relative to low-ability students, so the extent to which lower-income families benefit depends on the correlation between ability and family income on entry to secondary school.

Nechyba (2000) introduces private school vouchers into a public school system where there are heterogeneous housing types within each neighbourhood and thus incomplete stratification. Private schools are restricted in offering a single tuition price to all students, but they are able to choose whom to admit based on child's ability. He contrasts vouchers restricted to low income families or low income neighbourhoods with universal vouchers. In simulations he is able to show that residential stratification falls following the introduction of vouchers. Once again, the beneficiaries are high ability students, particularly those from poor households.

Ferreyra (2005) compares two choice systems to a neighbourhood schooling model. In the first, universal vouchers are issued that can be used in any private school. In the second, the vouchers are restricted to the non-sectarian (i.e. non-religious) schooling sector. She shows that both programs increase private school enrollment and affect household residential choice. However, under non-sectarian vouchers private school enrollment expands less than under universal vouchers. The effect of the vouchers on school quality and social welfare depends on the size of the voucher which is financed through taxation. The majority of the population benefits, slightly, from low value vouchers, but the proportion benefiting falls as voucher size increases. In simulations, the average winner is less wealthy yet more strongly prefers Catholic education than the average loser. Wealthy households, who already enjoy a high level of school quality before vouchers, tend to lose under high voucher levels regardless of their school choices due to their high income tax burden and capital losses in the housing market.

3.5.2 School choice in ACE models

An independent literature is being developed (mostly by non-economists) that models choice programs in an agent-based computational environment. These are non-equilibrium models, so changes in patterns of stratification over time as a program is introduced are of particular interest. Also, since they are entirely computational, spatial aspects of the models (such as where schools are located in relation to each other) are more explicitly considered.

Lauen (2003) uses an agent-based modelling approach to examine enrollment flows among eight public schools operating a choice program. Households live in a fixed location with a neighbourhood school, but are able to attend an alternative schools if there is spare capacity. They choose a school based on a utility function that includes distance of school to a fixed home address and the mean achievement level of the school. An important feature of the model is that schools are able to expand, contract and close, and the paper shows that the rules by which schools choose to change capacity have important effects on measures such as distance travelled, between-school variance in achievement, and between-school variance in enrollment. This dynamic of changes in school capacity on the supply side is interesting. However, the major drawback of this paper is that households are not able to move house because the price mechanism in the market for housing is not explicitly modelled.

Abernathy and Mackie (2002) construct a model of school choice between public (non-selective) schools and private schools that can select-out lower ability pupils. In each round of the simulation, households are able to apply to public and private schools with assignment based on chance (public schools) or ability of child (private schools). They are guaranteed a place at the neighbourhood school. The characteristics of this model appear similar to those of Nechyba and Ferreyra, but they are able to observe net gains in social welfare through the sorting of pupils over time. However, as with Lauen, the major drawback of this simulation is that they do not incorporate a price mechanism for housing that determines neighbourhood assignment in the public sector. Therefore, in the context of the English schooling market, it is not clear how interesting this model is.

One reason why hedonic markets for housing have not yet been modelled in an agent-based context is that they are very difficult to mimic realistically in a simulation. Economists computing general equilibrium employ a Walrasian auctioneer approach, with all households simultaneously able to bid on all houses to achieve a Nash equilibrium. In the real world, few households are actively seeking to move at any one time; bidding on several houses at once is complex because only one can be purchased; chains develop whereby a contract to exchange on one house is contingent on another contract completing; and precise locational characteristics cannot be known because they are contingent on the house moves of others. However, modelling this process is important because the general equilibrium approach necessarily ignores the 'stickiness' that results from the transaction process and cannot be used to assess how quickly markets are likely to change.

3.6 Predicting the effect of English school choice reforms

One purpose of the previous section was to show how few theoretical models of school choice programs currently exist in the literature. None appear to be capable to contributing significantly to the debate about the effects of UK-style choice programs on school and housing stratification. In England, local markets for secondary school places are highly heterogeneous and so it would be impossible to develop a stylised model that is suitable for making generalised claims about the consequences of choice in all areas. However, one key feature is that proximity to school or catchment areas play an important role in house location decisions of almost every parent in every region.

This section explores the distinctive characteristics of the English schooling market and discusses how they might be translated into a stylised model of household choice. This market is characterised as being one where most houses have a 'de facto' neighbourhood school, with a very high probability of acceptance at that school, but also choice appears to be possible without moving house. This means that households purchase a joint bundle of a home, plus a set of probabilities of acceptance at schools given the house location. At the time of the house purchase, the outcome of the subsequent school admissions process is highly uncertain.

The significant choices available to parents arise via spare places at adjacent neighbourhood schools, allocated by some proximity criteria or catchment area priority; places at Voluntary-Aided schools, assigned using a religious adherence measure; and places at a small number of neighbourhood schools, made available to non-residents on the basis of aptitude for a subject. So, a model of school choice would need to incorporate these three school types, with schools and households located on a grid since geographical position is such a key determinant of the probability of admission. The grid should contain some heterogeneity in housing type and little spare housing capacity to reflect the constraints present in most of England. In this type of simple model there is no mechanism by which schools can choose to change their capacity or admissions policies because no school decision-making function is explicitly modelled.

Households in the model would need at least three attributes:

- income, y;
- child's initial measurement attainment, b, that is increasing in y;
- religious adherence indicator, r, where it is also likely that $E[r|y]$ is increasing. This determines the probability of acceptance at the religious school and also the utility derived from a religious education.

At its most basic, the household utility function (including a constant disutility from making long journeys to school, d, and a utility derived from the religious education a

school offers, s) would be:

$$U(b, y, r) = U(y - p, a(\theta, b), l(h), s(r), d)$$

It might also be desirable to consider the effects of the following attributes:

- a taste for educational achievement, τ, that is increasing in y (and even in b);
- a disutility for making long journeys to school, μ, that is decreasing in y;
- an additional group who acquire the religious adherence indicator at a cost, ω, but derive no utility from religious education.

The additional attributes would produce:

$$U(b, y, r, \tau, \mu, \omega) = U(y - p, a(\theta, b, \tau), l(h), s(r, \omega), d(\mu))$$

When, how and if households optimise is crucial to deciding whether it is appropriate to model the English system in an equilibrium framework. This is because the house moving process is very costly in a variety of respects, as described earlier, and this means that households may appear to act in a myopic fashion with respect to house purchases. In practice, secondary school quality enters the household utility function at different times for different households: it could be as early as a first house purchase prior to having children for some couples and as late as the point where choice of school must be expressed on an official admissions form for others.

3.6.1 Decision-making under uncertainty

We require households to optimise – choose their location on the grid – without knowing for certain the outcome of the subsequent school assignment process. This uncertainty requires individuals to act on the basis of current fixed beliefs about the chances of various outcomes taking place. There are several sources of uncertainty in the school choice process for parents in England, but it is the first that is arguably the most unique and complex:

Market uncertainty means that everyone is fully certain about their own preferences and endowments, but is unsure about the supply-demand offer of other economic agents. Specifically in the case of schools, they are uncertain about which schools are likely to offer them a place under different location options, in part because this is itself contingent on the preferences and endowments of other households.

Event uncertainty derives from exogenous occurrences as they affect the household. An example would be the risk inherent in a child sitting an exam as a test of aptitude

for entry to a school.

Information uncertainty would include lack of knowledge about their own child's ability, school qualities, probabilities of gaining places at schools, and even imperfect knowledge about other parent's preferences. In short, it is any ignorance that individuals can attempt to overcome by 'informational' actions designed to generate or otherwise acquire new knowledge.

Formally, we can conceptualise each household as having a set of possible *acts* (x), which are the purchase of each house-type in every district. The *states of nature* (z) are all the combinations of being offered or not offered a place at each of the schools in the model. For each combination of x and z there is the *consequence* $c(x, z)$ of choosing to live in that house-type in that district given z combination of schools offer places to the child. The *consequence* is the utility the household derive from their most preferred school offer in z, given their location in x. In general, the expected utility derived from an *act* x is the utility (*consequence*) derived from each *state of nature* multiplied by the expected probability, π, of that *state of nature* occuring (with probabilities summing to 1). But in the case of school choice this expected utility is very complex because the probability of *states* occuring and the utility derived if they do occured are both entirely contingent on the *act* of the individual. The probabilities alter because some schools offer priority to neighbourhood residents and the utility derived from *states of nature* vary by x because distance to school enters the household utility function.

In order to compute this model it is necessary to have household estimates of the probabilities of *states of nature* occurring, given x. It would be meaningless for households to be able to compute $\pi[z|x]$ perfectly so we need approximations to make the model plausible. The most reasonable assumption is that households rely on the past year's school admissions information in some imperfect manner. They may, for example, know the proportion of applicants accepted to schools in the past year or may acquire information based on the secondary school destinations of other children in their neighbourhood or their child's primary school. However, all these sources of information are highly imperfect.

It becomes obvious under this type of allocation why all households will treat the neighbourhood school as a type of 'insurance' policy since they are guaranteed a place. We might believe that some households are more risk averse than others in this decision-making process. A risk-averse household would strictly prefer a certainty consequence to any risky prospect whose mathematical expectation of consequences equals that certainty. However, deciding what type of households should have this risk-averse attribute is difficult because we can only empirically measure the revealed, inseparable blend of their valuations of the consequences and their attitudes towards risk.

3.7 Discussion

Devising a more efficient and equitable system of school admissions for England is extremely complex. Developing a model, suitable for the English context, that is capable of predicting the effect of changing admissions policies on housing and school stratification would seem to be a crucial tool to make available to policy-makers. Through a review of existing models of parental choice of school in the literature, this chapter has argued that there still needs to be considerable progress made before we have models that are useful to policy-makers who wish to make equitable and cost-effective decisions regarding admissions regulation, transport subsidies, the employment of choice advisors, and so on.

The sociological literature is descriptively interesting but appears reluctant to move towards making predictive claims about specific policies relative to well-defined counterfactuals. By contrast, the economics literature is capable to making very specific predictive claims, but is not yet well-adapted to the English context. In any case, it is currently so reductionist in terms of its perspective on human behaviour that it is unlikely to be making accurate predictions about the effects of policies on stratification. Given the rich sociological descriptions of parental decision-making described at the start of the chapter, economic models might be enhanced by modelling human behaviour in a more complex fashion than the current optimisation problem of simple maximisation of peer quality given a budget constraint.

Progress in economic modelling can only take place via a more detailed description of the current institutional environment and apparent household behaviour within the English context. Providing some of this contextualisation in terms of the importance of the housing market and specific admissions criteria in determining school allocations is the purpose of the next two chapters. However, many questions will remain unanswered. For example, we need a detailed understanding of why schools and housing are currently so much more *integrated* than basic models of catchment area schooling predict. Given this relatively high social integration we should be cautious about making claims that choice programs must, *a priori*, lower segregation. It is not impossible that they set in place a dynamic whereby a previously stable and integrated system becomes more stratified. As part of the development of this literature, economists should remain open to the possibility that inertia in the housing market, myopia by households, path dependency and psychic attachments to locations all mean that out-of-equilibrium analysis, as used by agent-based computational models, is an important counterpart to general equilibrium models.

Chapter 4

Sources of school stratification I: The role of the housing market

The analysis in this chapter is based around a thought experiment:

> *What would happen to school compositions if we removed all pupils from their current secondary schools and reallocated them based strictly on a policy of proximity to school?*

The research question has not been devised to provide causal estimates as to sources of stratification, described in the previous chapter. Instead it is intended to provide a rich description of sorting in both the school and housing market. Pupil postcodes in NPD are used to simulate this thought experiment, showing the current role of the housing market in contributing to school segregation by indicating the amount of current school segregation that is attributable to residential segregation. It also provides an estimate of the extent to which pupils do not attend their proximity school, giving an upper bound on the amount of choice exercised by parents. The simulation therefore helps us understand how this pupil mobility affects the stratification of pupils across schools. Through additional simulations it also identifies areas where the post-1988 choice policies might have resulted in increasing school segregation. In this way it contributes to the school segregation literature on the effects of the Education Reform Act using a technique that avoids the problems associated with longitudinal measurements of school segregation, described in the next section.

Unfortunately, the simulated proximity allocation used in this study, while insightful in certain respects, is a poor proxy for the real world experiment for one principal reason: if we abandoned school choice and non-proximity admissions criteria in England today, we would expect some reallocation in the housing market as parents move house to attempt to achieve their desired choice of school. In other words, residential levels

of segregation are endogenous to the policies under examination in this study, but no estimates exist as to the size of the endogeneity bias on residential segregation in our data.[1] Specifically, it is not clear the extent to which school choice reduces the need for parents to locate close to a good school, given the outcome of the English school choice process is highly uncertain for a family (as discussed in Chapter 3). This means that when an association between a policy and the size of post-residential sorting is measured, we can go no further than infer that the policy has increased school segregation or reduced residential segregation (or both).

The chapter begins by describing the school segregation literature in order to document recent changes in segregation, motivate the choice of segregation index used in this chapter and discuss the advantages of using the free school meals measure in cross-sectional rather than longitudinal data. The chapter then goes on to describe the simulation method and present findings from the simulation.

4.1 Literature on changes in school segregation

The main body of literature that measures the effect of changes in school admissions on the sorting of pupils calculates segregation indices for local authorities and describes how these indices change over time. Until relatively recently in England, the only data collected on the demographic composition of schools was via the Annual Schools Census, where schools declared the number of pupils known to be eligible for free school meals (FSM). This was used as a proxy for social disadvantage to compare the distribution of pupils who were eligible for FSM, versus those who were not, across state schools.

4.1.1 Main findings

The first major research programme using large-scale longitudinal quantitative data measured secondary school segregation in the years 1989 to 1999. It claimed to show that, contrary to popular opinion, school segregation actually fell in the period immediately following the Act and has risen slightly since then (Gorard et al., 2002b, 2003). This finding, shown in Figure 4.1, prompted a fierce debate as to whether the large fall in segregation between 1991 and 1993 should be causally attributable to choice policies 'unlocking the iron cage' of neighbourhood catchment areas, or alternatively whether there were technical reasons for the fall in their measure. The two technical arguments

[1] A fifth of parents had taken account of catchment areas the last time they moved home and those who hadn't done so were more likely to take other special actions to further their application (Coldron et al., 2008).

regarding their choice of index and stability in the FSM measure are discussed here since they are directly relevant to methods chosen for this chapter.

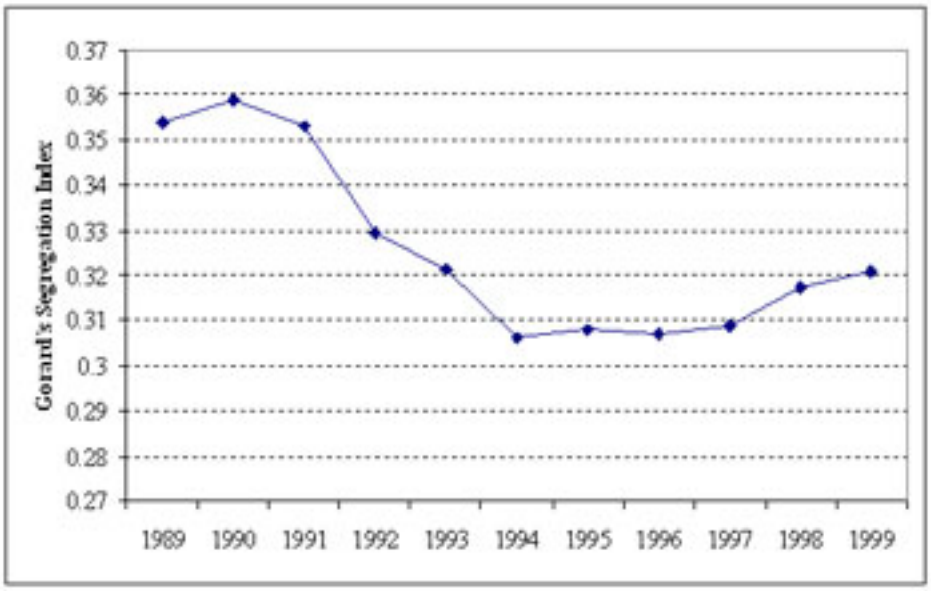

Figure 4.1: Values of Gorard's segregation index (GS)

All researchers, using very different measures of segregation, agree that secondary school segregation, measured using FSM, has been increasing very slowly from 1994 to 2007 (Allen and Vignoles, 2007; Department for Children, Schools and Families, 2008a; Goldstein and Noden, 2003; Noden, 2000). The increases in school segregation have been fastest in London (Allen and Vignoles, 2007) and in areas with high proportions of pupils in VA schools (Allen and Vignoles, 2007; Goldstein and Noden, 2003). However, these are measured associations and there is no suggestion that these are causal relationships.

4.1.2 Gorard's segregation index

The first technical argument relating to the validity of Gorard et al.'s findings relate to their choice of a non-standard segregation index. The formulae for Gorard's segregation index (GS) and the index of dissimilarity (D) for segregation in an LA with I schools and where FSM pupils are eligible for free school meals, $NONFSM$ pupils are not eligible for free school meals and there are N pupils in the LA are:

$$D = \frac{1}{2}\sum_{i=1}^{I}\left|\frac{fsm_i}{FSM} - \frac{nonfsm_i}{NONFSM}\right|$$

$$GS = \frac{1}{2}\sum_{i=1}^{I}\left|\frac{fsm_i}{FSM} - \frac{n_i}{N}\right|$$

Allen and Vignoles (2007) argue that the GS index is not the optimal way of measuring changes in school segregation because:

1. GS is not bounded by 0 and 1: the upper boundary varies according to FSM eligibility, making comparisons across LAs with varying levels of deprivation difficult. GS is better described as an 'indicator' rather than an index of segregation;

2. GS is not symmetric, meaning that it is capable of showing that FSM segregation is rising and NONFSM segregation is falling simultaneously; and

3. GS is actually systematically variant to changes in overall FSM eligibility, except in the most stringent and unlikely of circumstances. It has a tendency to fall as FSM eligibility rises, regardless of the change in the unevenness of school's shares of FSM and NONFSM pupils.

D and GS are closely related and measure similar aspects of unevenness segregation, but it is substantively important which index is selected. D and GS will only be highly correlated where levels of FSM eligibility do not vary greatly, and this is not the case across schools data. So, for example, Allen and Vignoles (2007) argue that GS overstates the magnitude of the fall in segregation after 1989 by 100 per cent, compared to D.

Noden (2000) and Goldstein and Noden (2003) also provided different criticisms of Gorard et al.'s work. However, the index of isolation, chosen by Noden, and the index that is implicit in the multilevel modelling approach are both also systematically variant to changes in overall FSM eligibility and are not 0-1 bounded. For example, the index of isolation has a tendency to rise as overall FSM eligibility rises, regardless of the change in the unevenness of schools' shares of FSM and NONFSM pupils. This chapter therefore uses the index of dissimilarity (D), which is a 0-1 bounded index, thus measuring the relative level of segregation compared to complete evenness and complete segregation.

4.1.3 Stability of FSM measure

The fall in all segregation indices in the early 1990s, illustrated in Figure 4.1, shows a pattern of two years of falling segregation up to 1993 followed by a rapid levelling out of the level of segregation. This particular pattern of data seems inconsistent with an explanation whereby a policy change produces a change in the sorting of pupils into secondary schools because secondary schools have five cohorts of pupils at any one time. Therefore, a one-off policy change should produce five years of falling

segregation (at roughly the same rate) as the cohorts of pupils who entered the school under the previous regime gradually reach the age of 16 and leave the Census.

The large fall in FSM segregation between schools between 1991 and 1993 coincided with a serious recession in England, and most researchers now agree that this was responsible for the changes in measured segregation between schools. The problem with the FSM measure is that FSM eligibility is not a fixed attribute of the child, but changes as the economic circumstances of the family change. So, it is possible that segregation can change from one year to the next as pupils switch their status from NONFSM to FSM, and vice versa, even if there is no change in the actual composition of the school. A recession would produce falling school segregation, for example, if the incidence of unemployment disproportionately affected families with children in relatively affluent schools. Recessions cannot, of course, affect the FSM status of pupils who are already eligible and this is why the *relative* deprivation of already deprived schools tends to fall.

If we accept the argument that the fall in school segregation in the early 1990s was due to the recession, it is also possible that the more recent rise in school segregation has simply been the product of an improving economic environment. Given the susceptibility of the FSM status indicator to changes in economic circumstances, Gibbons and Telhaj (2006) use a measure of KS2 attainment from NPD to show there was very little change in ability segregation between 1996 and 2002. However, ideally we also want a method for measuring FSM segregation without the problems described above. This chapter, by using a counterfactual in the same time period, overcomes the question of stability of the FSM measure.

4.2 Method and data

4.2.1 Data

The analysis in this chapter draws data from the National Pupil Database, Annual Schools Census and the 2001 Population and Housing Census of the UK, described more fully in Chapter 2. The 463,117 pupils in this analysis were in year 9 (age 13/14) in 2002/3. There is no postcode relating to year 6 or year 7 for this cohort in NPD, so there is a risk that the pupil has moved house since year 7 and the postcode used in this study will not relate to residential location at the point choice was made. Using a cohort in year 9 (rather than year 7) has one advantage that it means every student in England will have completed the secondary admissions process.[2]

[2] 7% of this year group transferred to secondary school at the end of year 7 or year 8

Key descriptors of all variables used in this study are in Table 4.1. Pupil social background and prior attainment are measured in four different ways in this chapter. First, social background is assessed using the FSM indicator. Second, high ability children are identified using an indicator for whether a child scored in the top quintile in KS2 tests at the end of primary schools. Third, a separate low ability indicator of whether the child scored in the lowest quintile in KS2 tests is used. This is relatively highly correlated with FSM. Finally, the child's total KS2 score is used as a continuous measure of ability on entry to secondary school.

Table 4.1: Summary statistics for LAs

	Weighted by pupil numbers			
Description	Mean	S.D.	Min	Max
Number of pupils in LA	5,139	3,559	361	13,157
Number of schools in LA	34	25	3	101
Average population density in LA	0	1	-0.77	5.82
% of pupils in LA at grammar schools	4%	9%	0%	42%
% of pupils in LA at VA comprehensives	14%	10%	0%	65%
% of pupils in LA at foundation comprehensives	16%	19%	0%	100%
% of pupils eligible for FSM	13%	9%	3%	64%
% of lone parent families	7%	1%	5%	14%
% of families with no parent in work	13%	6%	4%	49%
Mobility (pupils who move schools in simulation)	52%	14%	15%	86%
Median pupil's simulation journey as % of actual	64%	12%	31%	95%
FSM segregation (D) in real data	0.29	0.07	0.13	0.54
FSM segregation (D) in Simulation I	0.24	0.06	0.07	0.45
FSM post-residential sorting	0.05	0.05	-0.17	0.38
FSM post-residential sorting not due to grammars and VA ($seg^{real} - seg^{proxIII}$)	0.03	0.04	-0.11	0.20
Top 20% ability segregation (D) in real data	0.27	0.14	0.06	0.71
Top 20% ability segregation (D) in Simulation I	0.16	0.03	0.04	0.27
Top 20% ability post-residential sorting	0.11	0.14	-0.05	0.58
Top 20% ability post-residential sorting not due to grammars and VA ($seg^{real} - seg^{proxIII}$)	0.03	0.03	-0.05	0.28
Low 20% ability segregation (D) in real data	0.23	0.08	0.09	0.51
Low segregation (D) in Simulation I	0.16	0.04	0.01	0.29
Low post-residential sorting	0.07	0.07	-0.05	0.40
Low post-residential sorting not due to grammars and VA ($seg^{real} - seg^{proxIII}$)	0.03	0.03	-0.04	0.22
Ability (cont.) segregation (ICC) in real data	0.12	0.10	0.01	0.47
Ability (cont.) segregation (ICC) in Simulation I	0.04	0.02	0.00	0.13
Ability (cont.) post-residential sorting	0.09	0.10	-0.02	0.45
Ability (cont.) post-residential sorting not due to grammars and VA ($seg^{real} - seg^{proxIII}$)	0.02	0.02	-0.01	0.18

4.2.2 A counterfactual to current school admissions

The policies we label as 'school choice' policies in the UK are all intended to reduce the strength of the relationship between place of residence and school attended. Therefore, one possible counterfactual to these policies is the administrative allocation of pupils to school based solely on a proximity admissions criterion. This study tests two key propositions. First, that the amount of pupil mobility in an LA depends on the ability of parents to access a non-proximity school. This in turn is related to:[3]

- population density in the LA, which indicates the size of the potential choice set for parents;[4]
- the proportion of lone parent families and families where no one works in the household;
- the proportion of schools in the LA who may be using non-proximity oversubscription criteria, i.e. grammar, Voluntary-Aided (non-grammar) and Foundation (non-grammar) schools.

Second, that the difference in the level of current school segregation relative to segregation under the proximity allocation will be greatest where grammar, VA and Foundation schools exist. This may be because oversubscription criteria at these schools tend to favour higher income or ability pupils. Alternatively, this may be because school choice strategies differ by social background of family.

If this second proposition holds, we can then infer that the potential to reduce school stratification via a policy intervention that institutes a strict proximity policy is greatest in areas with many grammar, VA and Foundation schools. However, we must accept that residential sorting may rise in response; therefore gains in school integration may represent an upper bound.

The proximity counterfactual is created by computer using the OS grid references for pupil and school addresses in the dataset using the following rules:

1. all schools must fill their places on the basis of proximity, with a strict preference for pupils who live closer to the school. All other school admissions criteria at schools are removed (though single sex schools remain single sex);[5]
2. school capacity is set as either the current official school capacity or the size of

[3]Spare capacity in the LA was also tested but never significant, so is not reported here

[4]Population density is used as a proxy for the size of the parents choice set. Alternatives such as average number of competitor schools within a certain radius (Levačić, 2004) or drive-time (Burgess et al., 2007) are relatively highly correlated at LA level with average population density.

[5]The simulations retain single sex schools since gender is considered a fixed non-SES characteristic and so to avoid an upward bias on the effect of population density on mobility levels and a downward bias in the effect on sorting

the current pupil intake, whichever is greater;[6]

3. no parents will be allowed to exercise choice to attend another school, even if that school has spare capacity;[7]

The simulation presented in this report is essentially a Priority Matching Mechanism, one of several two-sided matching mechanisms described in Chapter 1.[8] The method in the simulation is as follows:

- The first round identifies each pupil's nearest school and distance from home and pupils are allocated to it, provided that there is enough spare capacity for them. If there are more pupils for whom it is the nearest school than there are places available, only those nearest are allocated.
- At the end of the first round, 84 per cent of pupils have been allocated to their nearest school. Forty-one per cent of schools (1,271 of 3,071) are full to capacity and will therefore be excluded from future iterations.
- The second round is similar, but with fewer spare places left at schools and only 16 per cent of pupils. The nearest school (with remaining space) is identified for each unplaced pupil. Again, pupils are allocated to the nearest school if there is enough space to accommodate them, with priority given in strict order of proximity.

The process must be repeated 9 times to allocate all pupils to a school. At the end of the process there are still 33,845 empty spaces at schools (7 per cent of all places available at the start of the allocation) because there is spare capacity in the system. Table 4.2 summarises the key details of distances travelled by pupils who are allocated at each stage of the procedure.

The simulation does not intend to replicate a real-world situation since this type of strict proximity allocation, without grammar schools and church schools, has never existed in England. Even where proximity is now the over-riding allocation principle in an area, many LAs continue to use catchment areas to aid planning and certainty of the allocation process for parents. It is simply suggested that this simulation provides

[6]Because of the need to exclude certain pupils from the allocation process, e.g. borders, the % capacity utilisation for 2000 is first calculated on the basis on total pupils at school. Where it is greater than 100%, it is reset to 100%. Simulation school capacity = pupils currently at school who are in simulation / % capacity utilisation

[7]The capacity constraint is significant in the simulation to the extent that if every pupil were allocated to their nearest school, 41% of schools would exceed their current capacity. Furthermore, if every school is filled to its current size but on the basis of proximity, 23% of pupils are allocated to more than one school and 27% of pupils are not allocated at all.

[8]Many alternative specifications were tested, but do not alter the substantive findings of this chapter. These specifications include allocation to nearest school without constraining capacity and an allocation mechanism whereby priority to oversubscribed schools is given to the pupils for whom the next nearest spare place is furthest from their home.

Table 4.2: Summary statistics for proximity allocation procedure

Round	Pupils needing places	Pupils placed in round	Minimum journey to allocated school (km)	Median journey to allocated school (km)	Maximum journey to allocated school (km)
1	463,117	388,311 (84%)	0.0	0.9	38.8
2	74,806	47,107 (63%)	0.2	2.4	25.5
3	27,699	16,690 (60%)	1.0	3.6	35.9
4	11,009	6,619 (60%)	1.3	5.2	43.8
5	4,390	2,798 (64%)	2.5	6.3	44.7
6	1,592	802 (50%)	4.1	7.4	46.3
7	790	464 (59%)	6.2	8.9	82.2
8	326	256 (79%)	8.3	13.3	15.8
9	70	70 (100%)	13.0	14.4	16.1

a valid mechanism for examining the stratification implications of current student sorting. This claim is made on the basis that only a tiny proportion of the 463,117 pupils are required to make unreasonable journeys in the allocation.[9]

Given that almost all grammar and VA schools existed prior to the 1980s, two additional simulations are used to provide a better indication of the possible role of new choice or post-1988 policies in changing pupil sorting. The three simulations are referred to as:

Proximity 1: the proximity allocation where no schools retain current pupils (this is the principal simulation and is described above)

Proximity 2: a proximity allocation where grammar schools retain current pupils (to generate an estimate of between-comprehensive pupil sorting)

Proximity 3: a proximity allocation where grammar and VA schools retain current pupils (to generate an estimate of between non-faith comprehensive pupil sorting)

4.2.3 Measuring differences in segregation

Since the outcome of interest in this study is the stratification of a local education market under different scenarios, it is necessary to geographically define the market. This study principally relies on LAs as the area of analysis, despite the high levels of cross-LA movement between certain LAs, particularly those in London. The simulation allows pupils to attend their nearest school, even if it is in a different LA to their

[9]For example, 449 pupils would be required to travel more than 5 times their current journey distance; 730 pupils would be required to travel over 10 km (though for some this is their nearest school)

home or current school. However, the measurement of LA-level segregation in the actual data versus the proximity allocation will often involve slightly different sets of pupils.[10]

Segregation between schools in an LA is measured using the index of dissimilarity (D). Segregation by FSM eligibility, by top 20% ability and by the lowest 20% ability is calculated using D. In addition the intra-class correlation (ICC) of a continuous measure of KS2 ability between-schools in an LA is used as a further estimate of segregation, where $ICC = 1$ indicates that schools are fully stratified because there is no within-school variation in ability and $ICC = 0$ indicates that schools are fully integrated by ability because there is no between-school variation in ability.

This study seeks to account for differences in the level of actual segregation between schools, seg^{real}, versus the level of segregation under the proximity allocation, seg^{proxI}. It is hypothesised that the size of this difference will relate to the level of pupil mobility (i.e. the proportion of pupils who are not currently educated in their proximity allocation school) in LA i, but also structural features of the quasi-market:

$$seg_i^{real} - seg_i^{proxI} = \beta_0 + \beta_1 mobility_i + \epsilon_i$$

$$seg_i^{real} - seg_i^{proxI} = \beta_0 + \beta_1 grammar_i + \beta_2 VA_i + \beta_3 foundation_i + \epsilon_i$$

Given the segregation index is 0-1 bounded, there is no clear a priori reason to favour the measurement of post-residential sorting as the absolute difference in the value of the segregation indices ($seg^{real} - seg^{proxI}$) over the proportionate difference in the value of the segregation indices (seg^{real}/seg^{proxI}). A problem with (seg^{real}/seg^{proxI}), particularly if D is used, is that it will have a very high value if an LA with very low levels of segregation experiences a small absolute increase in segregation that we might think has little effect on social welfare (e.g. the value of D doubles). It would be greater in this case, for example, than a higher segregation LA that experiences quite a large absolute, but smaller proportionate increase in segregation. The rank of LAs on both the absolute and proportionate measures of post-residential sorting will be sensitive to the segregation index chosen. So, both are tested and the absolute difference is selected on the basis that it provides more consistent regression results.

There are two properties of D that are highly relevant to its use as a dependent variable in a regression and therefore warrant mention here. First, D incorporates a linear payoff criterion to unevenness in the distribution of FSM pupils across schools (Zoloth, 1976). If we believe that the effects on social welfare of schools having different FSM

[10] Alternative specifications of all regressions were tested using 105 areas created by combining LAs where there are significant cross-LA movements, but did not alter the substantive findings. Results are available for this specification of the model.

proportions are non-linear, then D can provide us a broadly acceptable ordinal ranking of segregation but an inappropriate cardinal measurement of amounts of segregation. The robustness of results to this issue can be mitigated to some extent by replication of results using an index, such as the Square Root index (Hutchens, 2004), with a highly non-linear payoff criterion.[11]

A second issue is that the value of segregation under a random allocation of pupils to schools will be significantly greater than zero because a single school cohort is quite small, and this is generally an issue where a segregation index is used as the dependent variable in a regression (Carrington and Troske, 1997). This is known as the random allocation bias, the size of which is a function of the size of the LA, the number of schools in the LA and the overall FSM proportion in the LA (Cortese et al., 1976). Randomisation tests (presented in Allen et al., 2008) show that the variation in the size of the random allocation bias between LAs is substantial in all single cohort NPD datasets and therefore potentially invalidates all existing NPD segregation research that cannot account for this. The random allocation bias is overcome in the specifications described above because the size of the bias will be almost identical for seg^{real} and seg^{proxI} (since their margins should be approximately the same in the LA), thus removing most of the influence of the bias on the dependent variable $seg^{real} - seg^{proxI}$.

Because this chapter compares school segregation to a counterfactual proximity allocation of pupils, it is similar in its strengths and limitations to Burgess et al. (2007). They try to explain why the magnitude of the difference between residential and school segregation (post-residential sorting) differs substantially between LAs. However, because residential segregation is measured using wards as sub-units, which vary systematically and substantively in size across England, the random allocation bias is serious and confounds their measured association of post-residential sorting with population density.[12] This study moves beyond the scope of their work by analysing the role of own-admissions authority schools in contributing to sorting and also overcomes important methodological concerns regarding a random allocation bias, as described above, on their measure of post-residential sorting.

4.3 Results

4.3.1 The level of pupil mobility in LAs

The term mobility in this analysis means the extent to which pupils attend a different school to their current school in the proximity allocation. If the year 9 pupils in this

[11] These robustness tests are not reported in this chapter, but are available.

[12] Urban LAs tend to have more schools, relative to the number of wards, which increases their post-residential sorting ratio under a random allocation.

simulation were re-allocated to schools on the basis of strict proximity, 52% of pupils would be placed in a different school to their current school. This is consistent with the Burgess et al. (2007) indication that approximately half of all pupils currently attend their nearest school. The median pupil's distance to school under a proximity allocation would be just 64% of the length of their current journey.

Levels of pupil mobility differ substantially by LA, as shown in Figure 4.2. Not surprisingly, high mobility LAs are largely located in or around London, with 86% of pupils in Lambeth schools, for example, not attending the school they would be under the proximity allocation. Manchester, Liverpool and Birmingham also have high levels of pupil mobility. By contrast, areas where very little pupil mobility currently appears to be exercised are more rural, for example Leicestershire (mobility = 15%), East Riding of Yorkshire (23%), Rutland (24%) and Cambridgeshire (27%).

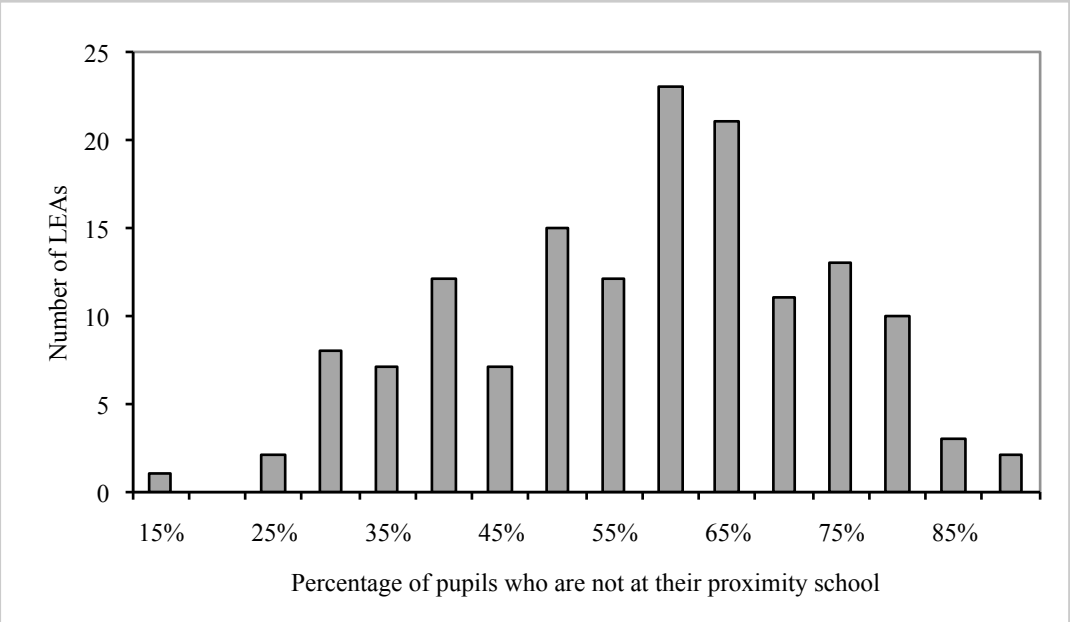

Figure 4.2: Distribution of mobility levels across LAs

Table 4.3 shows the estimates from the regression to test whether pupil mobility in an LA is associated with particular school or population characteristics. The overall model explains 75% of the between-LA variability in the proportion of pupils who are not at their proximity school; population density and all school type variables are significant at 1%. The size of the effect of grammar schools and VA schools is approximately the same, which is logical since neither tends give priority to pupils on the basis of proximity. Though the Foundation school coefficient is significant, the size of the effect on pupil mobility is very small: if 25% of pupils in an LA are at Foundation schools, mobility would be estimated to increase by just 2.5 percentage points. The effect of population density is so strong that an urban LA with only community comprehensives

might have higher pupil mobility than a very rural area with grammar schools.[13]

Table 4.3: Association between the quasi-market and level of mobility

Dependent variable	*mobility*		
Number of observations	147, weighted for no. of pupils in LA		
$Adj.R^2$	75%		
	Coeff.	t-stat	$P > \|t\|$
Proportion of pupils at grammar schools	0.6999	10.43	0.001
Proportion of pupils at VA schools	0.6147	8.31	0.001
Proportion of pupils at foundation schools	0.0987	30.06	0.003
Population density	0.0513	5.49	0.001
Proportion of lone parent families	1.1384	1.82	0.071
Proportion with no parent in work	0.0831	0.53	0.594
Constant	0.2927	7.04	0.001

Fifty-two percent of pupils do not attend their proximity school, yet Figure 4.3 uses additional simulations and analysis of the pupil's neighbourhood to establish that the destinations of only around one in five pupils might have been affected by post-1988 choice policies (this estimate ignores the possibility that changes in capacity at schools have altered the 'neighbourhood' school for some pupils). It estimates this because many of these sorting pupils are at grammar or VA schools or are likely to be at the de facto neighbourhood school, but this is not identified by the simulations.

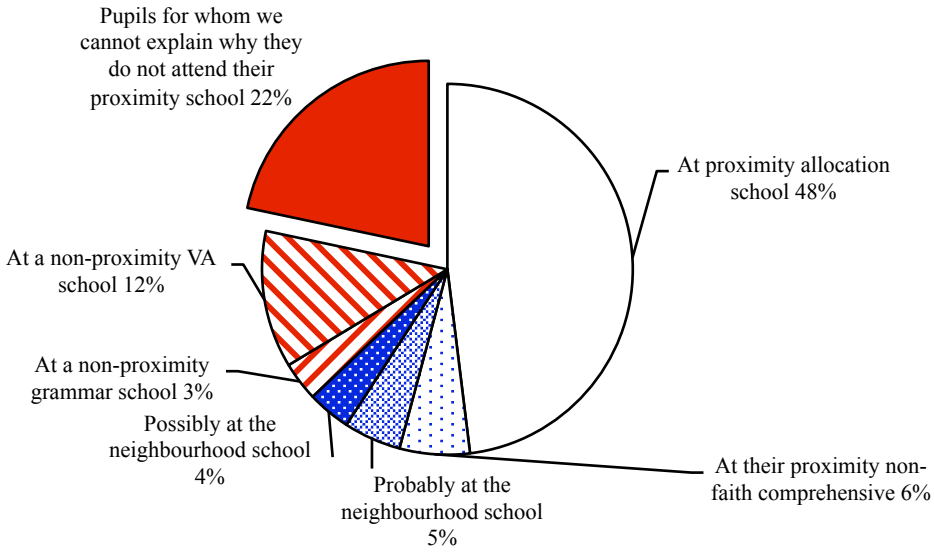

Figure 4.3: Understanding why pupils do not attend their proximity school

[13]These findings are very similar to those of Taylor (2001), who uses earlier data from six local authorities and finds that the proportion of parents who reject their nearest school is greatest in the more urban LAs.

Fifteen per cent of pupils in England are attending a grammar (3%) or VA school (12%), and this is not their proximity school. Six per cent of pupils attend their proximity non-faith comprehensive (or secondary modern) school, i.e. the proximity school in Proximity I is actually a grammar or VA school. Five per cent of pupils are attending the same school as the majority of pupils living in their Middle SOA (or neighbourhood). We can assert that these pupils are highly likely to be attending a designated catchment area school, or they are attending a school under a proximity criterion. An additional 4 per cent of pupils are probably at their school on the basis of catchment/proximity criterion because they attend the same school as the majority of pupils in the Lower SOA where they live. Since the Lower SOA is smaller than the Middle SOA, this identifies a set of pupils where the Middle SOA is crossing the catchment boundaries for more than one school. However, the use of the lower SOA (which is a small area so can have a homogenous social mix) presents a risk of mis-classification of pupils, so it might be the case that all the pupils in the SOA are successfully 'exercising choice' to attend a school far away from their homes.

This analysis shows us that more than half the pupils who are not attending their proximity school are either highly likely to be at a neighbourhood school or they are choosing not to based on criteria available to parents before 1988 (i.e. grammar and faith schools). Whereas Burgess et al. (2007) suggested that there is a high level of choice in England because 50% of pupils are not at their nearest school, this research places an upper bound on pupil mobility due to the new choice policies of 22%. This appears to be more in-line with the only existing study to model exact catchment areas, which found pupil mobility of around one-third in an LA with some grammar schools (Parsons et al., 2000). It is also consistent with a nationally representative survey where 25% of parents report they chose not to apply to their nearest maintained school in 2006 (Coldron et al., 2008).

Unfortunately, this analysis allows us to do no more than speculate about why the remaining 22% of pupils are not at their proximity school. These pupils are more likely to live in London (where they make up 34% of all London pupils) and are not, on average, attending schools with a superior social mix or ability of intake to their proximity allocation school. We can hypothesise that these 103,223 pupils are not at their proximity school for the following reasons:

1. **Edge of catchment.** Some pupils are attending the same school as pupils in the neighbourhood, but the SOA categories could not identify this (for example, because catchment boundaries cross the SOA). Half of these unaccounted for pupils are attending the same school as over 50% of pupils in the Output Area (but this statistical area is very small, so we cannot be certain that it is a neighbourhood school).

2. **Family relocation.** Some pupils may have gained their place at the school on the basis of a sibling policy, where other members of the family were allocated a proximity place because the family used to live the area. Alternatively, the family may have moved house since gaining a proximity place in Year 7.

3. **LA-wide ability banding.** Some pupils are attending a school on the basis of a banding system in an LA or school: this may, or may not be, their first choice school. The four London LAs who were operating LA banding at the time these pupils entered secondary school do have relatively high degrees of unexplained mobility.

4. **Choice without displacement.** Some pupils are exercising choice to attend a school other than the proximity school, and have gained a place at their chosen school via a proximity criterion (or other means) because all pupils who live closer either gained a place at the school or did not wish to attend it.

5. **Choice with displacement.** Some pupils have successfully exercised choice, having been allocated a place at the school via a non-proximity criterion such as aptitude (or as the sibling of a child who gained a place via an aptitude criterion), thereby displacing pupils who live nearer to the school than they do. There is some evidence for this phenomenon in the data: 23% of the unexplained mobility pupils are attending Foundation schools, versus 18% of the population and West et al. (2004) note that Foundation schools are more likely to have selective admissions criteria than community comprehensives.

6. **Displacement due to others' choice.** If we believe that choice with displacement explains a significant proportion of this unexplained sorting, then we should expect a significant number of pupils to be not attending their (non-faith comprehensive) proximity school because they are rejected in favour of pupils who live further away from the school than they do, but who gain places on non-proximity criteria.

4.3.2 School segregation and post-residential sorting

In this dataset, the weighted mean school segregation in LAs, measured by the dissimilarity index, is 0.29 by FSM and 0.27 by top 20% ability. The level of school segregation differs substantially by LA, as shown in Figure 4.4.

The regressions of LA level school segregation on a set of school type variables in Table 4.4 shows that higher levels of school segregation are associated with greater proportions of pupils educated at grammar and VA schools (this identifies the possible effect of these schools without endogenous residential sorting confounding estimates,

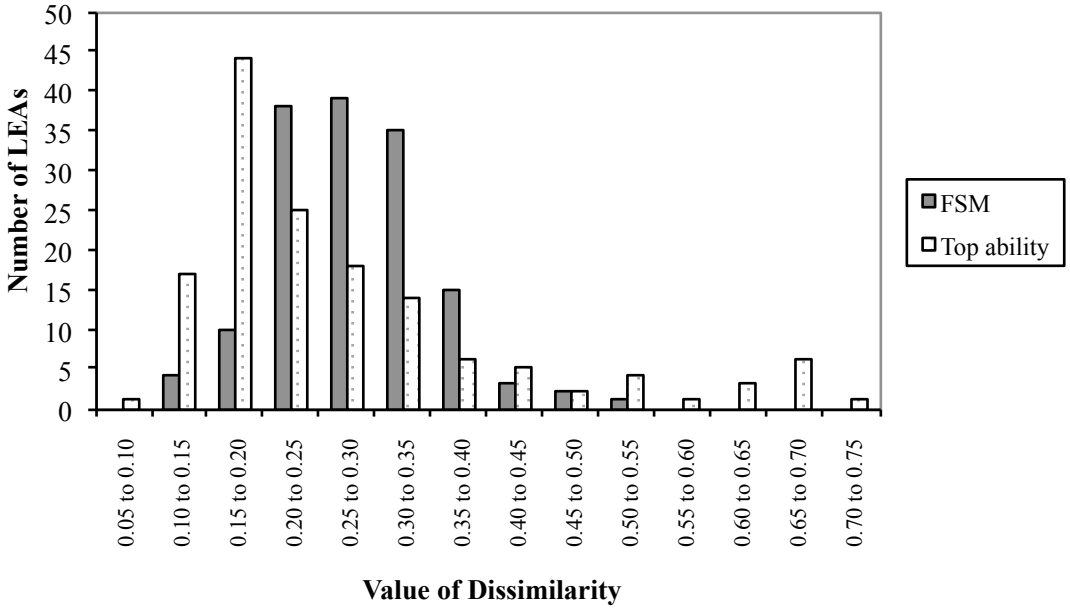

Figure 4.4: School segregation in English LAs

but is not a causal impact since the supply of school places is not explicitly modelled). Foundation schools are also associated with segregation by ability, but not FSM segregation in this sample. As with all regressions reported in this study, explanatory variables explain a relatively low proportion of variability in FSM segregation and a high proportion of variability in ability segregation. We cannot know whether this is due to the poorness of FSM as a proxy for social disadvantage, or whether school types have a clearer, more direct effect on ability sorting.

Table 4.4: Association between school segregation and school types in LEA

Dependent variable	seg^{real}			
Number of observations	147, weighted for no. of pupils in LA			
	FSM (D)	Top ability (D)	Low ability (D)	Ability ICC (D)
$Adj.R^2$	22%	86%	72%	89%
Proportion of pupils at grammar schools	0.320 (5.88) **	1.397 (28.57) **	0.696 (17.97) **	1.055 (32.94) **
Proportion of pupils at VA schools	0.206 (3.43) **	0.187 (3.46) **	0.203 (4.75) **	0.139 (3.94) **
Proportion of pupils at foundation schools	0.030 (1.13) n.s.	0.084 (3.56) **	0.048 (2.59) **	0.073 (4.74) **
Population density	-0.009 (-1.25) n.s.	0.003 (0.41) n.s.	0.001 (0.20) n.s.	0.002 (0.60) n.s.
Proportion of lone parent families	-0.080 (-0.16) n.s.	1.150 (2.52) *	0.505 (1.40) n.s.	0.420 (1.41) n.s.
Proportion with no parent in work	-0.091 (-0.72) *	-0.065 (-0.57) n.s.	-0.192 (-2.13) *	-0.050 (-0.67) n.s.
Constant	0.258 (7.65) **	0.101 (3.34) n.s.	0.155 (6.48) **	0.025 (1.24) n.s.

Almost all LAs have lower school segregation in the simulation proximity allocation of pupils to schools (seg^{proxI}), the values of which directly result from residential segregation. This is an important finding since it cannot easily be reconciled with the suggestion that low income families are the principal beneficiaries of policies that reduce the role of residential location in school admissions.

Figure 4.5 maps the distribution of seg^{real} and seg^{proxI} in LAs by FSM and top ability. The proximity segregation distribution represents the best possible reduction in school segregation by instituting a proximity policy. Overall, the weighted mean level of post-residential sorting is 0.05 (s.d. 0.04) by FSM and 0.11 (s.d. 0.12) by top ability. In both cases, segregation under the proximity allocation is typically lower. However, several LAs have a lower calculated segregation between schools currently than in the proximity allocation, i.e. post-residential sorting is less than zero. This is most likely to be because large numbers of pupils are crossing LA boundaries and so the calculation of segregation contains different pupils in each instance, i.e. the LA is not the valid market. Alternatively, whether by chance or design, these LAs have catchment areas drawn around schools that do not reflect strict proximity well and result in lower segregation.

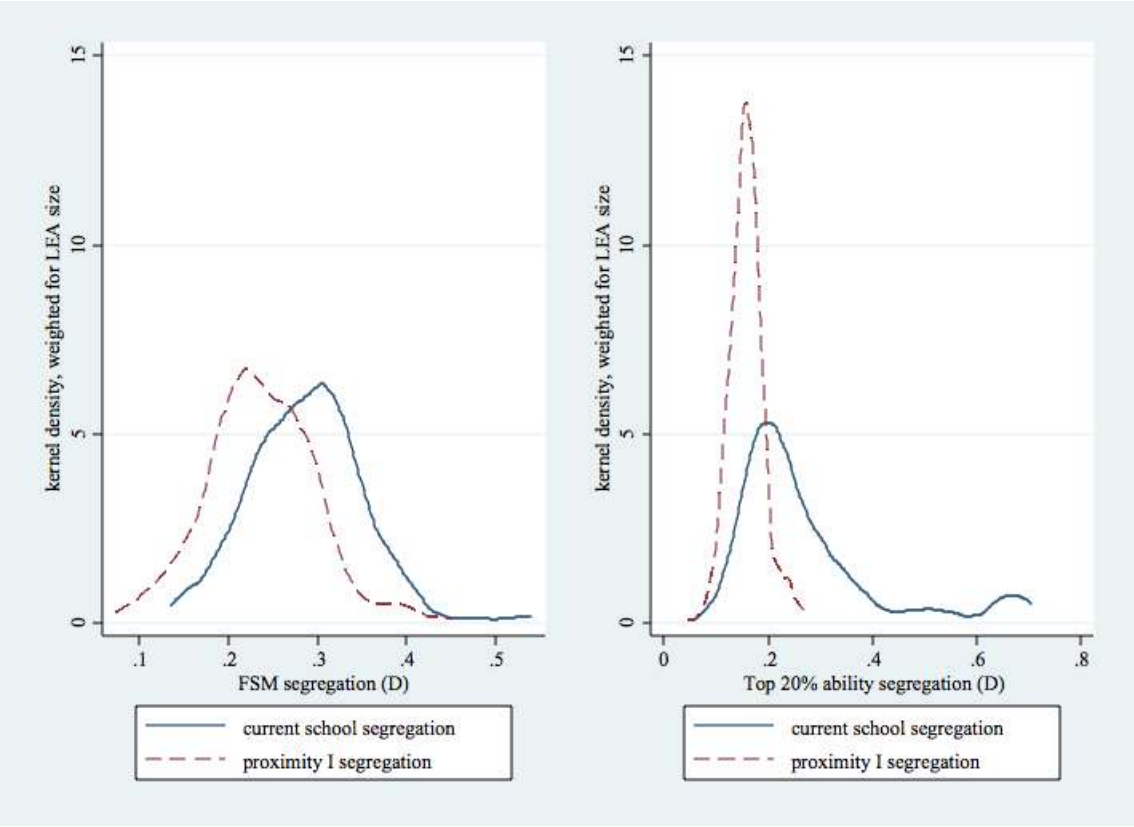

Figure 4.5: Current school versus Proximity I segregation

4.3.3 Explaining levels of post-residential sorting

There is a positive and significant correlation between levels of mobility in an LA and post-residential sorting by FSM and ability. These correlations, shown in Figure 4.6, are 0.48, 0.56 and 0.62 for FSM, top ability and low ability, respectively. The outliers by top ability post-residential sorting are areas with grammar schools.

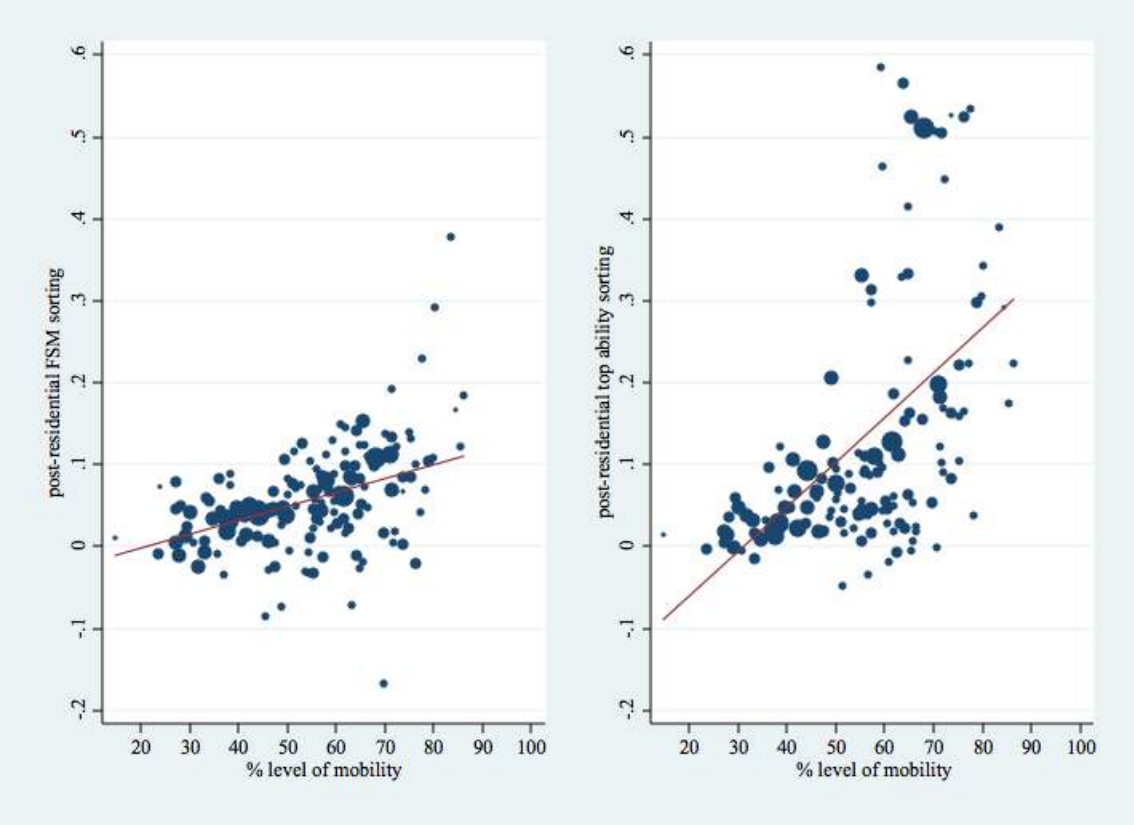

Figure 4.6: Relationship between pupil mobility and post-residential sorting

Table 4.5 shows that LAs with greater proportions of Voluntary-Aided, Foundation or grammar schools have higher levels of post-residential ability and FSM sorting. Not surprisingly, the effect of grammar schools on top ability sorting is very high indeed: where grammar schools in an LA educate 25% of the pupils, they contribute to a 0.35 unit increase in top ability post-residential sorting. VA and Foundation schools control their own admissions, yet their typical effect on school sorting is different. VA non-grammar schools contribute to higher levels of post-residential sorting than Foundation non-grammar schools. All these findings are robust to the exclusion of London LAs from the regressions. There is relatively weak evidence that post-residential sorting is lower in areas of high unemployment.

Table 4.5: Results from post-residential sorting regression

Dependent variable	$seg^{real} - seg^{prox}$			
Number of observations	147, weighted for no. of pupils in LA			
	FSM (D)	Top ability (D)	Low ability (D)	Ability ICC (D)
$Adj.R^2$	33%	90%	82%	92%
Proportion of pupils at grammar schools	0.215 (5.56) **	1.400 (33.35) **	0.686 (23.29) **	1.059 (39.14) **
Proportion of pupils at VA schools	0.099 (2.34) *	0.159 (3.43) **	0.162 (4.98) **	0.133 (4.47) **
Proportion of pupils at foundation schools	0.045 (2.42) *	0.073 (3.62) **	0.035 (2.50) **	0.067 (5.16) **
Population density	0.015 (2.76) **	0.024 (4.05) **	0.021 (5.12) **	0.012 (3.14) **
Proportion of lone parent families	-0.026 (-0.07) n.s.	0.827 (2.11) *	0.171 (0.62) n.s.	0.282 (1.12) n.s.
Proportion with no parent in work	-0.004 (-0.04) n.s.	-0.240 (-2.46) *	-0.208 (-3.04) **	-0.094 (-1.50) n.s.
Constant	0.022 (0.92) n.s.	-0.005 (-0.22) n.s.	0.033 (1.79) n.s.	0.005 (0.29) n.s.

As pupils succeed in attending a non-proximity school this mobility raises school segregation relative to residential (or the proximity) segregation. We can use the proximity allocation to show that pupils who are not eligible for FSM are more likely to benefit from current sorting than FSM pupils, where 'benefit' is defined to mean they improve their peer group relative to the proximity allocation. Pupils who remain in the same school following the simulation may not have mobility, but this does not mean they do not benefit from sorting. We can say they will benefit if the choice policies produce an improved peer group at the school they currently attend.

One way to illustrate the peer group of a school is as the proportion of pupils eligible for FSM in the real data and Proximity I. Table 4.6 shows the change in the FSM peer group experienced by FSM and NONFSM pupils as a result of current mobility. Sixty-one per cent of FSM pupils are worse off in terms of their peer group under current sorting, compared to a proximity allocation. Just 13 per cent of FSM pupils have mobility that results in them improving their peer group and 25 per cent find that their own school's peer group improves through current mobility. By contrast, half of the pupils not eligible for FSM have a better peer group under current sorting (and some have a considerably better peer group, which is why more than 50 per cent of pupils are worse off in the proximity allocation overall).

Table 4.6: Gains from mobility as measured by change in FSM peer group

	Worse peer group than under proximity:		Better peer group than under proximity:	
	% in a different school	% remaining in the same school	% in a different school	% remaining in the same school
Pupils eligible for FSM	29.86	31.45	13.30	25.38
Pupils not eligible for FSM	29.16	20.86	19.62	30.36

School segregation is said to matter because it creates inequalities in the school peer group experienced by children from different backgrounds. Standard estimates of the size of the peer group effect suggest a one s.d. improvement in a child's peer group leads to around a 0.1 s.d. change in that child's achievement (see Ammermueller and Pischke, 2006, for a recent summary). The proximity simulations suggest that the typical child currently experiencing a relatively poor peer group (one s.d. below the mean) would see this improve somewhat under a proximity allocation to around 0.57 s.d. below the mean peer group. Though this is a meaningful improvement in peer group and therefore predictive educational outcomes for the child, it would be wrong to claim that neighbourhood schooling offers a magic bullet for lowering social class inequalities in achievement.

4.3.4 Post-1988 mobility and post-residential sorting

The prediction of sociologists working in the field of education has been that the policies introduced since 1988 will increase pupil mobility and that this will in turn produce greater school segregation (e.g. Gewirtz et al., 1995). Two simulations – Proximity II and III – can separate mobility likely to have existed pre-1988 from potentially newer mobility. Proximity III allows grammar and VA schools to keep existing pupils (to identify the post-residential sorting attributable to these schools) and Proximity II separates the contribution of VA schools to post-residential sorting by allowing only grammar schools to keep existing pupils.

Figure 4.7 shows that mobility caused by grammar and VA schools can explain most of the post-residential sorting of pupils by ability, but they explain little of the total FSM post-residential sorting.

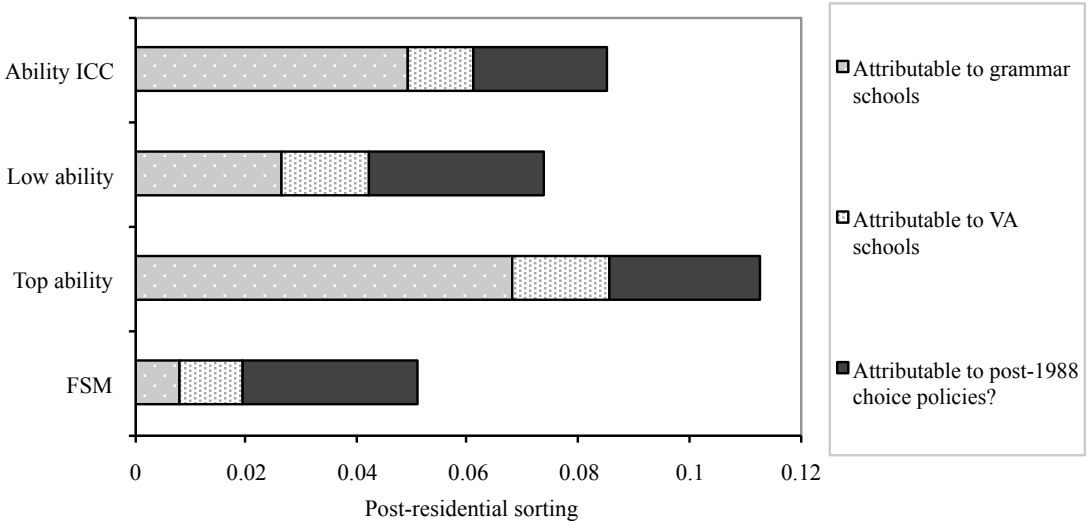

Figure 4.7: Sources of post-residential sorting

The weighted mean value of post-residential sorting that might be attributable to the post-1988 choice policies is 0.03 for each of FSM, top ability and low ability segregation. This is quite a low figure, but it varies between areas: it is as high as 0.20 in Haringey (by FSM); and 0.28 in Wandsworth (by high ability).

4.4 Discussion

This chapter has sought to show that we can improve our understanding of the impact of school choice policies on student sorting via simulations that reallocate pupils to secondary schools strictly on the basis of proximity by exploiting the availability of

pupil postcodes. This is a new (and therefore imperfect) technique and could be improved, for example: via consideration of each school's relevant geographical 'market'; by accounting directly for partial ability selection by comprehensives in the simulation; through repetition over several years to understand the changing role of the housing market; and by accounting for the part of the endogenous residential sorting that takes place between the ages of 5 and 11 as more years of pupil-level data become available.

Though half the pupils in the study do not actually attend their proximity allocation school, much of this can be likely attributed to the shape of catchment areas and the presence of grammar and VA schools (i.e. pre-1988 mobility). The simulations estimate that mobility between non-faith comprehensive schools is likely to involve no more than one-in-five pupils in England.

Pupil mobility results from the combination of both the choice to attend a non-proximity school and the displacement of local pupils from their proximity school and, if successfully exercised, naturally means longer journeys to school. The proximity allocation indicates that the typical journey currently made by a pupil is 60% longer than the minimum necessary. In fact, over 5 million kilometres[14] of additional travel is made by 11-16 year olds every school day, either because parents are choosing not to send their child to the local school, or because the local school is choosing not to give the child a place. This raises important efficiency issues that are little discussed in the literature. We should not assume that a parent incorporates the external costs of pollution and congestion, *even* if they rationally decide that a longer journey (with the associated time and money cost for parent and child) is warranted by the 'superior' education at the end of the journey. Furthermore, if school places are highly constrained and proximity is *not* the sole allocation rule, one parent's rational choice to access a 'superior' education farther away from home may force a pupil local to the school to make a longer journey to an 'inferior' one.

The simulations show that FSM and ability segregation is almost always lower in the proximity counterfactual than in the actual data, confirming that where pupils are sorting themselves into a non-proximity school, it does tend to increase social and ability segregation between schools, relative to underlying residential segregation. This would appear to be contrary to the suggestion that choice policies are likely to disproportionately benefit low income families because they were previously unable to afford homes close to popular schools, thus lowering school segregation (e.g. Hoxby, 2003a; Burgess et al., 2005b). In England we did not start from a position of complete residential stratification, and in addition we have given some schools both the means

[14] 1.4 million km total distance to school in real data versus 0.9 million km in proximity allocation. Difference of 0.5 million km is multiplied by 2 to incorporate the return journey and multiplied by 5 for the 5 year groups

and the motivation to recruit pupils with above average ability and those with fewer problems. However, because existing school allocations widely use catchment areas rather than proximity criteria, it is possible that these are also leading to the observed differences between school and residential segregation.

Unfortunately, we cannot use these findings to predict the size of the effect that removing discretion over admissions from grammar, VA and Foundation schools and forcing a strict proximity allocation would have on school segregation because this model cannot identify the magnitude of the endogeneity of residential sorting. We can say that a proximity allocation has the potential to lower school segregation by the amount that is post-residential sorting, and the size of this potential reduction is greatest in areas with grammar, VA and Foundation schools and higher population density areas. However, we do not know the extent to which residential segregation will rise in any one area to offset this potential fall, so post-residential sorting should be interpreted as the maximum possible reduction resulting from a proximity allocation. It is noteable, though, that areas with grammar, VA and Foundation schools do not currently have lower residential segregation than areas where all schools operate catchment areas.[15]

About two-thirds of LAs have a level of post-residential FSM sorting that is lower than 0.05, once segregation directly attributable to grammar and VA schools is accounted for. Though the simulations do measure significant mobility between non-faith comprehensives in these particular LAs, it would be wiser to attribute this to the inability of the simulation to capture the shape of traditional catchment areas for schools than it would be to use this to suggest that these parts of the UK are indisputably achieving choice without greater school segregation. Indeed post-residential FSM sorting that might be associated with post-1988 policies is only a weighted mean average 0.03 across LAs. This very low figure might explain why time-series of school segregation reveal so little change over time on a national basis. If policy-makers genuinely wish to improve equality of educational opportunity in England, this chapter suggests they should look closer at the continuing role of grammar schools, Voluntary-Aided faith schools and, most of all, of the proximity oversubscription criterion in producing socially segregated secondary schools.

[15] Also, despite the introduction of school choice, residential patterns (at least the ranking of LAs from lowest to highest) of economic segregation have remained fairly constant over the last twenty years (Meen, 2005). Some argue that residential segregation by income has actually intensified in London and other metropolitan areas in the UK over the past 20-30 years (Meen and Meen, 2003).

Chapter 5

Sources of school stratification II: School intake versus neighbourhood characteristics

Chapter 4 showed that local authorities with high proportions of schools that control their own admissions tend to have a large gap between the level of school and residential segregation. The purpose of this chapter is to document the extent to which these own-admissions schools have intakes that differ substantially from their local neighbourhood, and also to explore the role of specific admissions criteria in contributing to school segregation. Through the use of a survey of school admissions policies, a distinction can be drawn between the effects of own-admission schools that use explicit criteria designed to prioritise pupils with particular characteristics, and own-admissions schools who do not appear to have policies designed to cream-skim, but nevertheless do administer their admissions process.

One aim of the chapter is to distinguish between the phenomenon of schools cream-skimming pupils and stratification attributable to parental choice. This distinction is difficult to establish in the absence of data containing lists of applicants to schools, so the conclusions of this chapter are necessarily rather tentative. In fact, even with this type of data there are circumstances where this distinction would not be appropriate. For example, a school might discourage certain families from applying to a school through the use of complex opaque admissions procedures, including supplementary forms that ask for information on marital status, and so on. Thus, data on applicants to the school might appear to show that parental choice was responsible for the school's advantaged intake, yet school admissions procedures mean that true parental preferences were not reflected in applicant lists.

Two other studies have compared school and neighbourhood characteristics for secondary schools. The Sutton Trust (2006) compared the intakes of the 'top' 200 comprehensive schools (according the GCSE league tables) to their local neighbourhoods. It found evidence of social selection in that only 5.6 per cent of pupils were eligible for FSM in these schools, compared to 11.5 percent living in their local communities. Chamberlain et al. (2006) compared the intakes of all secondary schools to their postcode district, showing the Voluntary-Aided schools admit fewer FSM and more high ability pupils than are in their local neighbourhood. This chapter makes two additional contributions to evidence. First, it makes use of different measures of a school's 'neighbourhood', which it argues are superior to the use of the postcode district. Second, it is able to relate statistics on intakes directly to data to school admissions policies, therefore moving the debate beyond a discussion about the effect of governance on cream-skimming and towards the more nuanced debate about the use of specific admissions criteria and cream-skimming.

The chapter is structured as follows. First, the current intake characteristics of schools by school type are described. Second, the method of comparison with the local neighbourhood – the proximity ratio and recruitment ratio – are described and evaluated in relation to each other. Third, admissions policy data are used to assess how policies are correlated with overt cream-skimming. Fourth, the data on Voluntary-Aided schools is interrogated further to make preliminary suggestions as to why these schools tend to have advantaged intakes.

5.1 Schools included in analysis

This chapter uses the National Pupil Database of school leavers in 2005 for England, details of which can be found in Chapter 2. By examining the characteristics of school leavers, the intention of this chapter is to investigate the total effects of a school's admissions procedures, including the effects of admissions and exclusions after year 7. However, it does mean that the neighbourhood characteristics are measured with reference to the pupils postcode at 16 and not at the point they enter the school.

The purpose of this chapter is to examine the admissions policies and practices of two distinct types of non-grammar schools that control their admissions: those that have a religious denomination and those that do not. This distinction is made because criteria to establish religious adherence are complex, and would appear to focus on the characteristics of the family rather than the child. LA controlled Community schools are used as the comparison group for both these types of schools. There are currently four types of own-admissions schools: Academies, City Technology Colleges (CTCs), Foundation schools and Voluntary-Aided (VA) schools. However, because Academies

are still very new and CTCs are few in number, this chapter compares Community schools to (i) all Foundation schools (most of which will be non-religious former Grant-Maintained schools); and (ii) VA schools with a religious denomination. There are 42 VA schools that are not defined as religious, and these are excluded from the analysis because it is the religious admissions criteria that are of particular interest here. On the other hand, there are a very small number of Voluntary-Controlled schools that are defined as religious, but these are excluded from the *Religious* schools group in this chapter because their admissions are decided and controlled by the LA. Thus, the label *Religious* in this analysis refers only to VA religious schools. Table 5.1 marks in bold the three sets of schools that are compared in this chapter.

Table 5.1: Governance of religious and non-religious schools

	Academy	CTCs	Comm-unity	Found-ation	Vol. Aided	Vol. Con.	Total
Non-religious	10	14	1961	492	42	55	2574
Religious	4	0	0	9	**480**	38	531
Total	14	14	**1961**	**501**	522	93	3105

Since this research is concerned with the effect of admissions policies on school composition, the 164 grammar schools are excluded from the majority of this analysis. Table 5.2 shows that a large proportion of grammar schools have Foundation status, but since all grammar schools use very similar academic entrance tests, any intake difference between the 77 Foundation and 41 Community grammars would not seem likely to inform the analysis of 'comprehensive' schools that follows. Although grammar schools are excluded from the analysis, pupils who are currently in grammar schools are not dropped from the dataset, but instead are included when the neighbourhood characteristics of the schools in the analysis are calculated. This means that the average school in the analysis will have an intake that is less able than the neighbourhood in which it is located.

Table 5.2: Governance of grammar schools

	Non-grammar schools	Grammar schools	All schools
Not in analysis	135	28	163
Community	1920	41	1961
Foundation	424	77	501
Voluntary-aided religious	462	18	480
All schools	2941	164	3105

5.2 Current intake characteristics by school type

Table 5.3 shows the social, ethnic and ability characteristics of pupils who attend Community, Foundation and Religious schools, compared to pupils in all state schools. Foundation schools have the fewest free-school meals (FSM) eligible pupils (an average of 9.7%), followed by Religious schools (12.7%) and Community schools (15.1%). By contrast, Religious schools have the greatest proportion of pupils who scored in the top 25% nationally in Key Stage Two tests (28.3%), compared to Foundation schools (23.6%) and Community schools (21.0%). Throughout the remainder of this chapter, 'top ability' is short-hand for these pupils scoring in the top quarter of the population in KS2 tests, i.e. the pupils deemed to be of highest academic ability on entry to secondary school. Similarly, 'low ability' is short-hand for pupils scoring in the lowest quarter of the population in KS2 tests.

Table 5.3: Characteristics of pupils in schools by analysis group (%)

	Including grammars				Excluding grammars			
	All	Community	Foundation	Religious	All	Community	Foundation	Religious
FSM	13.3	14.9	8.8	12.4	13.8	15.1	9.7	12.7
Top 25% ability	24.9	21.9	29.9	29.9	22.7	21.0	23.6	28.3
Low 25% ability	25.2	27.8	21.4	19.0	26.2	28.3	24.0	19.6
Asian Indian	2.3	2.3	2.8	1.7	2.2	2.3	2.7	1.5
Asian Pakistani	2.3	2.8	1.5	0.9	2.3	2.8	1.5	0.8
Asian Bangladeshi	0.9	1.1	0.4	0.3	0.9	1.1	0.5	0.3
Black African	1.6	1.3	1.1	3.3	1.6	1.3	1.2	3.4
Black Caribbean	1.5	1.3	1.2	2.7	1.5	1.3	1.3	2.8
White British	81.5	81.9	82.4	77.2	81.5	81.9	82.4	77.5

Religious school have greater numbers of Black African and Caribbean pupils, but little reference is made to ethnic composition of schools for the remainder of this chapter. Though Community schools are clearly less affluent and have a lower ability profile than Foundation and Religious schools, the averages mask considerable diversity within these three sectors. Figure 5.1 shows that there are schools in all three sectors with FSM proportions above 50%, as there are schools with few FSM pupils.

5.3 Neighbourhood characteristics by school type

The chapter now turns to comparisons of a school's current pupil composition with the characteristics of the neighbourhood within which the school is located. There are

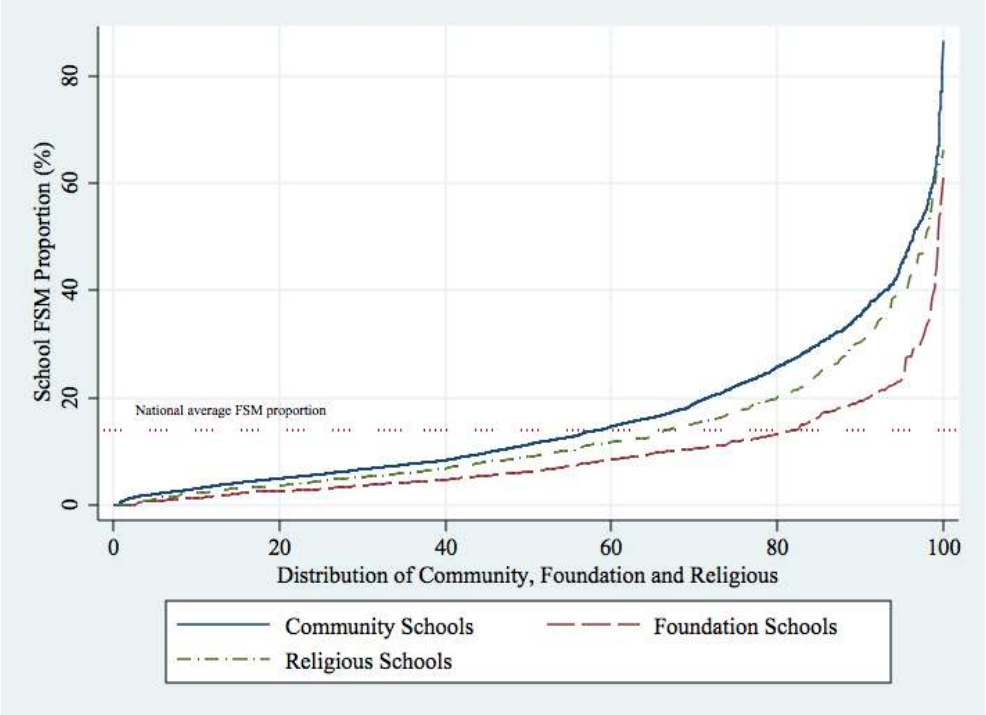

Figure 5.1: Distribution of FSM pupils within Community, Foundation and Religious sectors

many different ways we can conceptualise the neighbourhood, for example:

- the street adjacent to the school;
- a 3 kilometre walk zone for the school;
- the postcode area or district in which the school is located;
- the local authority in which the school is located.

There are problems with all definitions of the neighbourhood. For example, a 3 kilometre walk zone would be a good definition for a very urban school, but would exclude most pupils who attend their local rural secondary school. Local authorities (LAs) differ substantially in size across England so that shire county LAs are far too large to be considered a neighbourhood, yet London LAs are inappropriate as the neighbourhood for schools who are located on LA borders. The decision to use postcode areas (the first half of the postcode) by Chamberlain et al. (2006) and Sutton Trust (2006) to measure the characteristics of a school's neighbourhood means that, particularly in urban areas, the definition of neighbourhood is far broader than the recruitment area of the school. In rural areas the reverse is true because a Community school is likely to recruit from several villages, each of which may be in a unique postcode area. The use of postcode areas is also unlikely to reflect the neighbourhood for all schools located on the edge of a postcode sector. The choice of neighbourhood measurements reported

here are in some respects also necessarily arbitrary and flawed, but many have been tested and the substantive results reported are robust to the use of many alternative measures of neighbourhood.

5.3.1 Proximity and recruitment neighbourhoods

The first measurement of the neighbourhood characteristics used in this chapter is called the *proximity neighbourhood*. The conceptual idea is that it measures the characteristics of the school if all pupils in the neighbourhood wished to attend the school and the school allocated places simply based on proximity to the school. In computational terms, the number of pupils currently at the school is counted, then the computer finds that number of pupils who live closest to the schools from all pupils currently in state schools.[1] The social and ability characteristics of these *proximity* pupils are measured.

Table 5.4 shows the *proximity neighbourhood* characteristics for the school types in this analysis. Foundation schools are located in more affluent areas than Community schools, on average (neighbourhood FSM proportions of 9.6% versus 12.4%); Religious schools are in more deprived areas than both (neighbourhood FSM proportion of 16.5%). The proportion of top ability pupils living in the neighbourhoods for the three school types is more evenly balanced.

Table 5.4: Proximity characteristics of schools in analysis (%)

	Excluding grammars		
	Community	Foundation	Religious
FSM	12.4	9.6	16.5
Top ability	25.0	27.2	25.3
Low ability	24.8	22.8	25.2
Asian Indian	2.1	2.4	2.9
Asian Pakistani	1.9	1.4	2.9
Asian Bangladeshi	0.9	0.5	1.5
Black African	0.8	0.8	1.9
Black Caribbean	1.1	1.1	2.1
White British	84.6	83.5	77.7

The second measurement of the neighbourhood characteristics used in this chapter is called the *recruitment neighbourhood*. The conceptual idea is that it identifies the characteristics of all pupils who live within a similar distance of the school to three-quarters of the current intake (i.e. the school's main recruiting area). In computational terms, the distance of each pupil currently at the school is calculated and the 75^{th} percentile

[1] For single sex schools, only pupils of the relevant sex are included in the computation. However, for other types of school, including religious schools, all pupils are included.

home-school distance is recorded. Then all children who are living closer than the 75th percentile home-school distance are found, and their social and ability characteristics are measured. Table 5.5 shows the *recruitment neighbourhood* characteristics of schools in the analysis. The statistics are very similar to those in Table 5.4.

Table 5.5: Recruitment characteristics of schools in analysis (%)

	Excluding grammars		
	Community	Foundation	Religious
FSM	13.1	10.0	17.1
Top ability	24.5	26.4	23.9
Low ability	25.4	23.6	26.3
Asian Indian	2.2	2.4	3.0
Asian Pakistani	2.2	1.7	3.0
Asian Bangladeshi	1.0	0.5	1.4
Black African	0.8	0.9	1.9
Black Caribbean	1.2	1.3	2.2
White British	83.9	83.0	77.6

The choice of the 75th, rather than the 95th or 55th, percentile radius appears quite arbitrary. It represents a compromise between wanting a radius that captures the journey lengths for the majority of pupils at the school, yet excludes pupils who appear to make unrealistically long journeys. Table 5.6 shows that the pupil at the 95th percentile can make a journey as far as 172 km to school each day! Some of these very high figures are likely to be caused by data errors, so reducing the length of the radius for the analysis avoids these problems.

Table 5.6: Home-school distance (in km) for nth percentile pupil in school

	Mean	S.D.	Min	Max
15th percentile radius	0.93	0.72	0.14	17.12
35th percentile radius	1.78	1.43	0.36	25.60
55th percentile radius	2.86	2.39	0.54	41.83
75th percentile radius	4.33	3.81	0.73	85.22
95th percentile radius	8.27	7.82	1.43	172.16

These two measures of neighbourhood characteristics produce indicators of a set of pupils who live within a particular radius of the school, with the *proximity neighbourhood* being tighter than the *recruitment neighbourhood*. In practice, actual neighbourhoods the school draws from are often irregular rather than circular in shape, particularly where geographical barriers such as motorways, hills and lakes affect journey times. This means a school might draw from its direct neighbourhood, as it conceptualises it, yet may have an intake with characteristics significantly different from the neighbourhood measured by these statistics. However, if the school's actual intake is

irregular due either to the school themselves selecting a favourable catchment area, or due to parental choice, it is desirable to use the circular neighbourhood because it abstracts from these processes. For the purposes of this analysis, all that matters is that the statistic measures the average neighbourhood characteristics for any particular school type without bias, and there is no reason to suppose this should not be so.

Neither measure of neighbourhood is, *a priori*, more superior in urban versus rural areas. The differences in a school's calculated neighbourhood will be most pronounced in areas with high levels of pupils sorting. However, in practice, the measured neighbourhood characteristics differ very little between measurement based on *proximity neighbourhood* versus the *recruitment neighbourhood*, as shown by the correlation coefficients between the two for key pupil characteristics (see Table 5.7).

Table 5.7: Correlation between alternative measures of neighbourhood characteristics

	Correlation coefficient
FSM proportion	0.91
Top ability proportion	0.85
Low ability proportion	0.84
Asian Indian proportion	0.91
Asian Pakistani proportion	0.86
Asian Bangladeshi proportion	0.92
Black African proportion	0.93
Black Caribbean proportion	0.92
White British proportion	0.95

There is a question as to how meaningful quite a narrow definition of a neighbourhood is when we know that families move close to schools in order to secure a place, i.e. the neighbourhood characteristics are necessarily endogenous to school type. This is more true for the *proximity neighbourhood* than it is for the much wider *recruitment neighbourhood*. It means that the interpretation of data in this chapter is necessarily quite narrow – we can simply investigate the characteristics of a school compared to a situation where it solely admitted the pupils *currently* living close to the school. Of course, if a school that is currently partially selective, for example, abandons this policy, this is likely to produce a change in the school's neighbourhood characteristics as well as its school composition.

There is a related issue regarding Religious schools, especially Roman Catholic (RC) schools that have never had a stated commitment to educate the local community. As this is the case, it could be asked why the neighbourhood characteristics would be of any interest for RC schools. However, given that these schools do recruit from a wide geographical area, the *recruitment neighbourhood* is also measuring the characteristics

of the pupils over a similarly wide area. It is true that the measure does not identify which children are Catholic within this area. Without knowing the Religious affiliation of all pupils in the dataset, or alternatively knowing who wished to attend an RC school but did not gain a place, we have few methods for assessing whether these schools are cream-skimming. This means that differences in neighbourhood versus school characteristics for RC schools will reflect (i) differences in the characteristics of Catholic versus non-Catholic families; (ii) application choices made by families to these schools; and (iii) acceptance rates at RC schools for different groups of families. The final section of this chapter gives an example of how this data restriction of not knowing religion of family can be overcome by looking at sorting within the group of pupils who attend RC schools in London, and who are therefore likely to have sufficiently proved their RC religious adherence.

5.4 School intake versus neighbourhood characteristics

The chapter now turns to visual comparisons of pupil characteristics for each school type compared to neighbourhood characteristics. Figure 5.2 presents a scatter diagram with each point representing a Religious school's FSM proportion versus its *proximity neighbourhood* characteristics (the chart for the *recruitment neighbourhood* looks very similar). A school such as school (a) is located in an area of high deprivation with over 60% FSM eligibility in the neighbourhood. The school itself has an approximately equal number of FSM pupils as the neighbourhood has, and thus takes its 'fair share' of such pupils. School (b) is also located in an area of very high deprivation with around half of all local children eligible for FSM. However, less than 5% of the pupils actually attending the school are eligible for FSM. By contrast, school (c) is in a relatively typical area of England with few FSM pupils in the neighbourhood, yet it has a FSM proportion of almost 40% at the school. In general, the pattern for Religious schools is that the school composition is more affluent than the neighbourhood the school is located in, so most points on the chart are below the 45 degree line (67%, versus 37% for Community schools). There are strong regional patterns though, with London Religious schools being, on average, far more affluent than their neighbourhood (shown in Figure 5.3).

Figure 5.4 compares the intake of Religious schools to the local neighbourhood in terms of proportion of pupils scoring in the top ability on KS2 tests. The neighbourhood here is measured as pupils within the *recruitment neighbourhood*, although the pattern is similar for the *proximity neighbourhood*. It shows that no Religious schools are situated in a neighbourhood where over 40% of pupils are classified as top ability pupils, yet 13% of Religious schools have over 40% top ability proportion in their actual intake. Overall, 62% of Religious schools have an intake with more top ability pupils than the

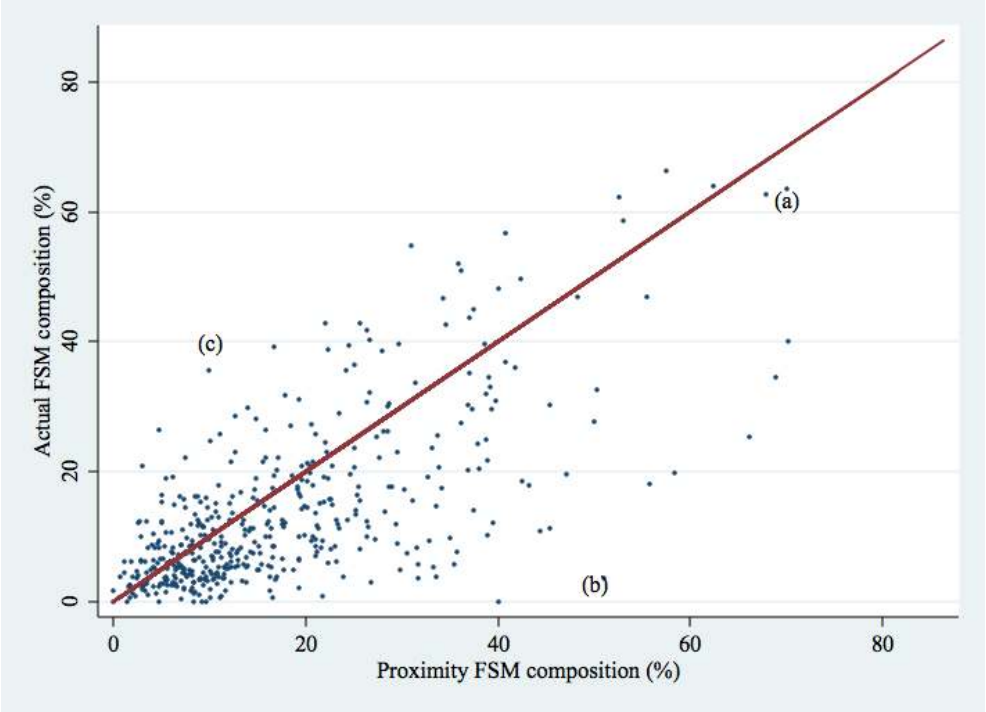

Figure 5.2: Actual versus proximity FSM composition for Religious schools

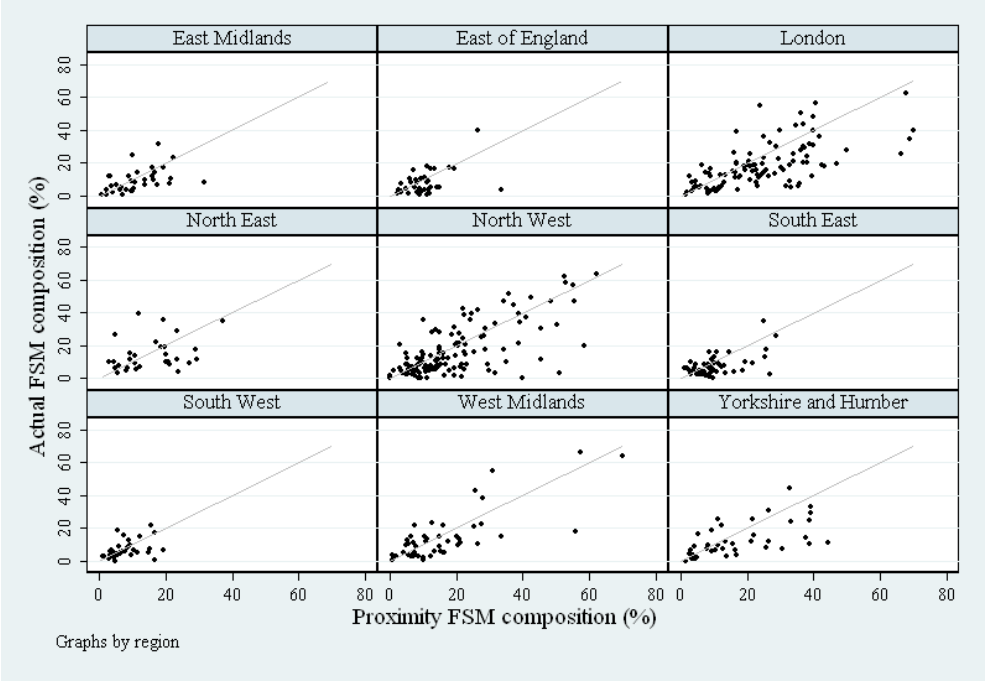

Figure 5.3: Regional differences in FSM composition for Religious schools

local neighbourhood.

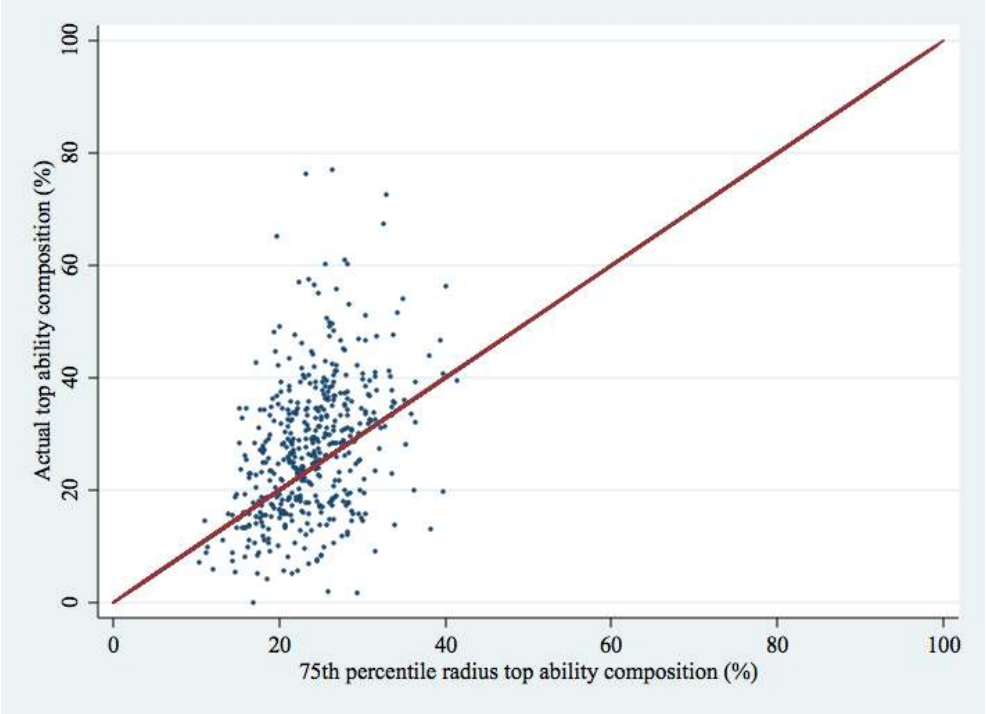

Figure 5.4: Actual versus proximity top ability composition for Religious schools

Figures 5.5 and 5.6 show the same charts for Foundation schools. Here the pattern is much less pronounced. Foundation schools tend not to be located in deprived areas, so schools such as that marked (a) on Figure 5.5 (a school with a deprived intake in a deprived neighbourhood) are unusual. Overall, 60% of Foundation schools (versus 37% Community schools) have a FSM proportion that is lower than their neighbourhood.

5.4.1 Proximity ratio and recruitment ratio

The graphical illustrations of neighbourhood and school intakes can be summarised in a ratio. Taking the example of FSM proportions in the school and neighbourhood, the *proximity ratio* is calculated as:

$$Proximity\ ratio = \frac{\%FSM_{school}}{\%FSM_{proximity}}$$

Table 5.8 shows the mean values of these ratios for FSM, ability and ethnicity characteristics of pupils at Community, Foundation and Religious comprehensive schools. The ratios using the *recruitment neighbourhood* for the denominator are shown in square brackets. For FSM ratios, a value greater than 1 indicates the school's intake

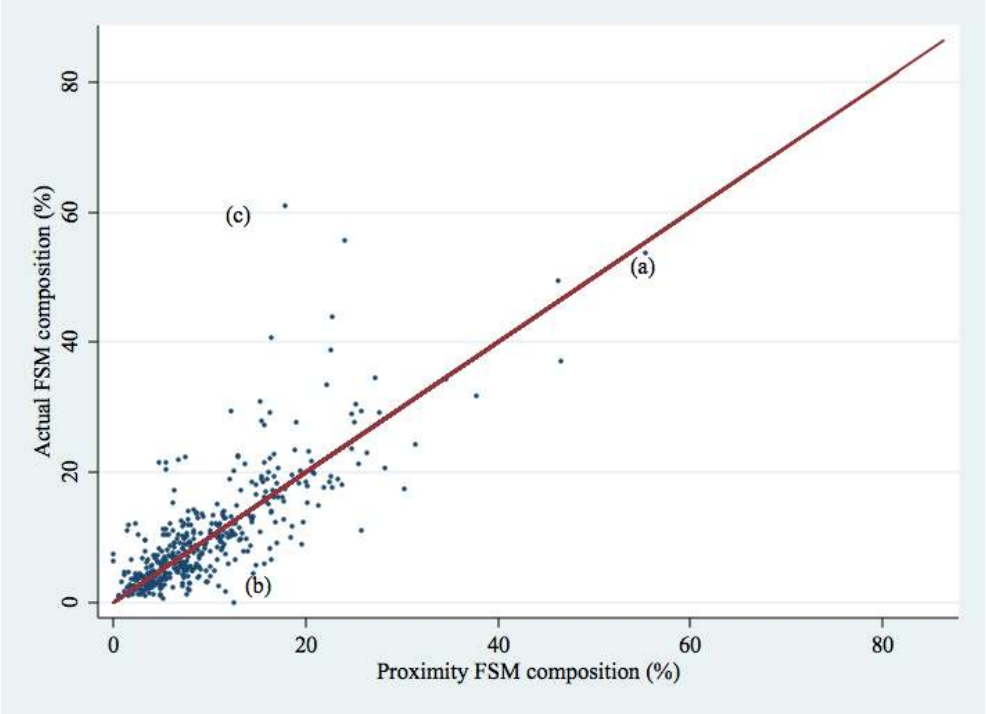

Figure 5.5: Actual versus proximity FSM composition for Foundation schools

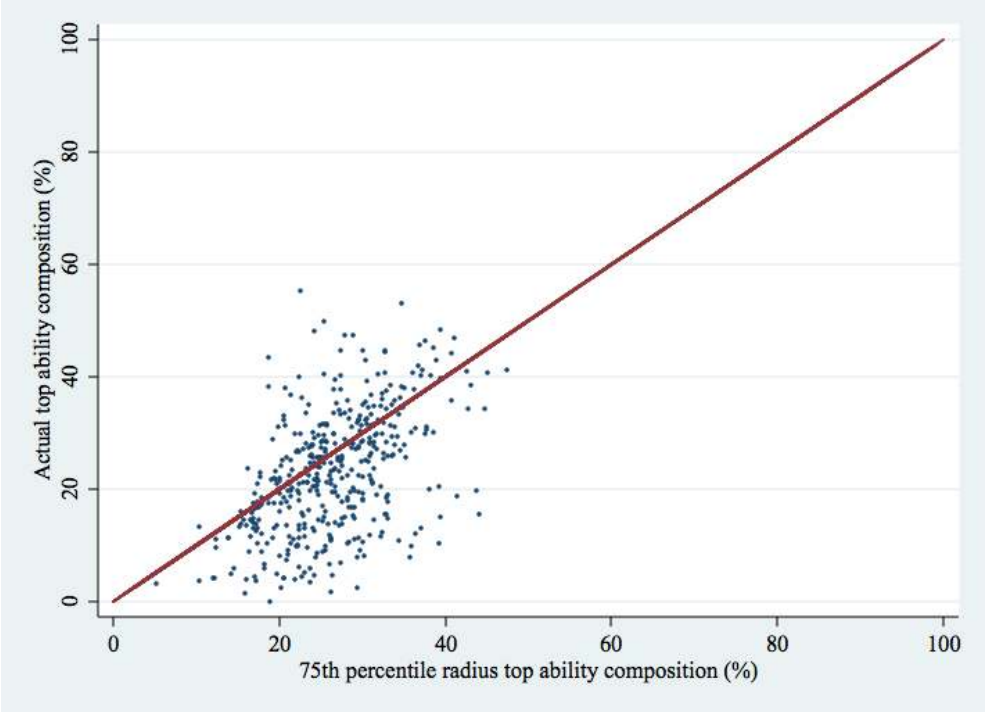

Figure 5.6: Actual versus proximity top ability composition for Foundation schools

is more deprived than its neighbourhood. Community schools have approximately 30% more FSM pupils than live in the neighbourhood immediately adjacent to their schools. By contrast, Religious schools have almost 10% fewer FSM pupils than live in the local neighbourhood. It is important to point out once again that, while grammar schools, Academies and CTCs are excluded from the summaries of these ratios, pupils in these other schools are included when calculating the neighbourhood characteristics of schools. This means that the comprehensive schools shown here will, overall, have intakes that are slightly more deprived (with fewer top ability pupils) than their neighbourhoods. The correct comparison is therefore between different types of comprehensive schools.

Table 5.8: Proximity ratios (excluding grammar schools)

	Community	Foundation	Religious
FSM	1.30 [1.12]	1.11 [0.97]	0.93 [0.72]
Top ability	0.84 [0.85]	0.88 [0.90]	1.24 [1.19]
Low ability	1.16 [1.11]	1.08 [1.01]	0.84 [0.74]
Asian Indian	1.05 [1.00]	1.04 [0.97]	0.86 [0.92]
Asian Pakistani	1.24 [1.16]	1.08 [0.99]	0.71 [0.53]
Asian Bangladeshi	1.08 [1.08]	0.96 [0.94]	0.67 [0.42]
Black African	1.22 [1.54]	1.14 [1.35]	1.32 [2.38]
Black Caribbean	1.04 [1.03]	1.04 [0.94]	1.26 [1.14]
White British	0.95 [0.95]	0.97 [0.97]	1.06 [0.99]

Note: 75^{th} percentile radius ratios are in square brackets

Overall, both Religious and Foundation schools achieve intakes that are more affluent and more able than Community schools, once the characteristics of the local neighbourhoods are taken into account. The differences between Religious and Community schools are very large: the statistics suggest that a typical Religious school is likely to have around 50% more top ability pupils than a Community school located in the same type of neighbourhood. By contrast, a Foundation school in a similar position would have around 5% more top ability pupils than a Community school. Pair-wise t-test comparisons of Community schools with Religious and Foundation schools show that the differences between these school types are almost all statistically significant at the 5% level. The exceptions are non-significant ethnic composition differences between Community and Foundation schools for Asian Indian, Black African and Black Caribbean groups. There are also non-signficant ethnic composition differences between Community and Religious schools for Asian Indian pupils.

5.5 School admissions policies and sources of stratification

It is impossible to use administrative data to conclusively establish why schools have intakes that are markedly different from their local neighbourhood. In particular, we would like to distinguish between the two main processes by which this could take place. Taking for example, a school with fewer FSM pupils than is representative of their local neighbourhood, this might have resulted from FSM-eligible families not applying to the school or alternatively, they may have applied, but not gained a place due to the admissions criteria used by the school. We can call the latter phenomena 'cream-skimming', meaning a set of admissions criteria and processes that have the effect of systematically favouring higher ability or more affluent pupils. However, it is also possible for schools to *discourage* certain applicants from applying to the school, for example by using complex admissions procedures or by only marketing the school to affluent neighbourhoods. These types of activities by schools should also be included in the definition of 'cream-skimming' activities.

West et al. (2004) document the types of admissions criteria used by schools at the time the pupils in this dataset entered secondary school. They highlight 'potentially selective admissions criteria' that may be used by schools to favour higher ability and easier to teach pupils, thus excluding others. Table 5.9 shows the proportions of Community, Foundation and Religious schools in the dataset that are classified as using these potentially selective admissions criteria. It shows that these criteria are overwhelmingly used by Religious schools, mostly for the purposes of establishing religious adherence of the family. A significant number of Foundation schools also use ability or aptitude tests, even though they are deemed comprehensive schools.

Turning to a comparison of the admissions criteria used by schools and their proximity ratios, a simple count of the number of potentially selective criteria mentioned for each school is calculated. Values range from zero to seven. Figure 5.7 charts the average FSM proximity ratio by school type for schools using different numbers of potentially selective criteria. Perhaps not surprisingly, the general pattern seen is that the more potentially selective criteria are used, the lower the FSM proximity ratio. Indeed, the FSM proximity ratio is identical for Community, Foundation and Religious schools that have no potentially selective criteria. Furthermore, where three potentially selective criteria are used, the proximity ratio is little different between school types.

Figure 5.8 shows a similar pattern for the top ability proximity ratio. Community, Foundation and Religious schools that do not use potentially selective criteria have similar ratio values. The very high values of the top ability proximity ratio relate to Foundation and Religious schools that use four or more potentially selective criteria.

Table 5.9: Use of potentially selective admissions criteria in 2001

	Community %	Foundation %	Religious %
Sibling's academic record	0	0	0
Primary school report	0	1	1
Family connections	1	2	5
Religious criteria	0	0	87
Interview with parent	0	0	11
Interview with pupil	0	1	16
School ethos	0	1	11
Banding administered by school	0	2	6
Uneven banding applied by school	0	0	2
Test administered by school	1	7	5
Ability/aptitude selection	1	13	7
Other selective criteria	3	8	82

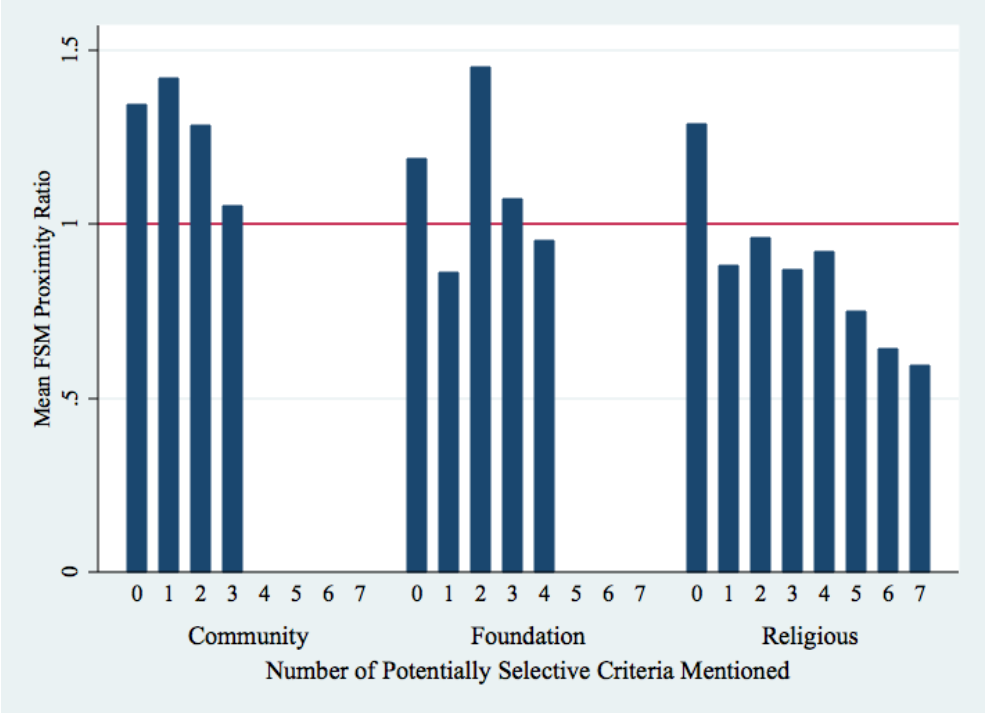

Figure 5.7: Mean FSM proximity ratio for schools grouped by admissions criteria

This analysis is quite crude because it classifies all potentially selective criteria equally, whereas some are more likely to skew the intake of a school than others. However, it does serve to show that the use of explicit potentially selective admissions criteria are an important source of stratification in the schooling system.

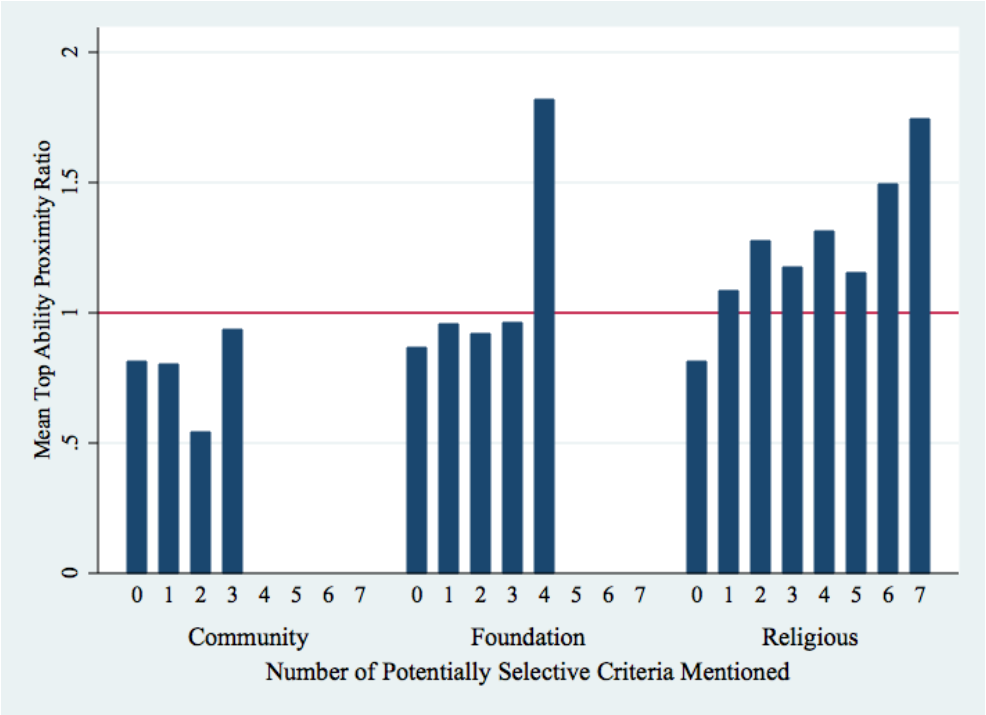

Figure 5.8: Mean top ability proximity ratio for schools grouped by admissions criteria

From this data it is notable that only a minority of Foundation schools have adopted admissions policies that are likely, by design, to raise the ability distribution of their intake. So, it does not appear to be true that schools who are given the means to cream-skim will necessarily take-up that opportunity. The interesting question is why some of these former Grant-Maintained schools chose to take the opportunity adopt selective admissions policies in the 1990s whereas others did not. The list of Foundation schools using 'potentially selective admissions criteria' shows that they tend to be located in education markets where other schools were selecting based on ability, prior to the 1990s. For example, they include many 'de-facto' secondary modern schools in selective schooling areas and comprehensive schools in local authorities where there are one or more stand-alone selective grammar schools. It is possible that their own change in admissions policy was motivated by the perception that the system of local school admission was already unfair, or that they felt it necessary to change their admissions policy in order to survive in a competitive environment.

5.6 Sorting within the religious schooling sector

This thesis returns to the causes of advantaged intakes at Foundation schools in Chapter 7, since they are almost all former Grant-Maintained schools. However, the statistics in the previous section showed that it is the Religious comprehensive schools that have, on average, the most advantaged intakes, so this section focuses exclusively on exploring this phenomenon more closely. The analysis is restricted to London since this is the region with sufficient Roman Catholic (RC) and Church of England (CofE) schools within close geographic proximity to each other, thus providing an environment where sorting within the Religious schooling sector is possible.[2].

While it is true that the average Religious secondary school has a more able and less deprived intake than their proximate neighbourhood, within the sector there are a few schools for whom these figures are exceptional (called 'élite' schools in this section). There are only two ways any school can achieve an intake that is far more affluent than their neighbourhood: either poor families living close to the school are choosing not to apply for a place (perhaps in expectation of rejection or because they do not find the school attractive); or these families are applying, but they are excluded by the school's admissions criteria. Clearly, the moral implications of a religious school excluding poor children are less ambiguous than if poor families do not wish their child to attend in the first place.

It is difficult to disentangle these competing choice versus cream-skimming hypotheses, but this section tries to contribute to this debate by looking at whether there are pupils who live close to, but do not attend, an élite Roman Catholic school, yet have a clear revealed preference for Catholic schooling because they attend another Catholic school. Given that the majority of pupils in Catholic secondary schools in London are required to be Roman Catholics (Pennell et al., 2007), these pupils will have satisfied Catholic credentials at another school. This section is interested in whether these non-élite-school pupils of RC parents are more likely to be of low ability or income status. In other words, do the five élite RC schools analysed here appear to cream-skim from within the pool of pupils from local Catholic families?

For this analysis, pupils who live 'close to' an élite RC school is defined in two ways:

1. the group of proximity children who live directly adjacent to the school and would attend if the school accepted only its closest pupils;
2. the (larger) group of children who live within a 3 kilometre radius of the school (i.e. a walk zone).

[2]The analysis in this section was carried out with advice from Anne West for Allen and West (2009)

Table 5.10 presents this analysis for five élite RC schools selected on the basis that their FSM *proximity ratio* was the lowest (0.25 or lower). The general pattern across these schools is that there are pupils who live close to these élite schools but who attend other RC schools. These non-attending RC pupils are more likely to be low ability or FSM eligible than the pupils who succeed in attending the élite school. Very few of these non-attending RC pupils are educated at a school with a higher ability peer group than the élite RC school they live near. This means it is unlikely that these non-attending RC pupils are choosing to attend another Catholic school which is both further away from their house and has a lower ability intake.

The 25 (female) proximity pupils who live very close to **School A** but attend an alternative Catholic school have almost four times the rate of the FSM eligibility as those who attend school A (28% versus 8%). They are also far more likely to have scored in the bottom quarter in KS2 tests.

School B also has a large number (41) of (male) pupils living very close to the school but attending alternative Catholic schools. These pupils have over four times the rate of FSM eligibility as pupils at school B (27% versus 6%). Most (but not all) of these pupils are attending another RC school with a lower ability peer group than the élite RC school.

Schools C and D are located in the same outer London LA. Just four pupils living very close to school C attend an alternative RC school; many more live within 3km and these pupils have much higher FSM eligibility rates. School D has more RC pupils living locally but not attending the school. Some of these are FSM eligible, but others are successfully attending school C, which has the high performing peer group.

School E has no alternative RC schools nearby, so few pupils live close to school E but attend another RC school. Those that do are not FSM eligible. This analysis of sorting within the Catholic sector cannot draw conclusions about the cream-skimming behaviour of this particular school. It does have an intake that is highly advantaged compared to its neighbourhood, but it is possible that this solely results from the socio-economic circumstances of Catholic families living in the area. Alternatively, there may be poorer Catholic families who wish to attend this school but do not gain a place. Because there are no alternative RC schools in the area, these children would be in non-religious schools, and so it is not possible to identify them as Catholic in NPD.

In all five schools, criteria and practices that could enable pupils to be 'selected in' or 'selected out' were used. According to the 2001 survey of school admissions, all five schools interviewed pupils and/or parents, one used banding and one used the academic record of siblings as part of the admissions process to decide which pupils of those who applied should be offered a place at the school.

Table 5.10: Cream-skimming analysis of 5 élite RC secondary schools

	N	FSM	Low ability	High ability peers?
School A = Inner London RC Girls School				
Current intake	143	8%	6%	
Proximity pupils who attend other RC schools	25	28%	13%	0%
3km radius pupils who attend other RC schools	101	24%	17%	0%
School B = Inner London RC Boys School				
Current intake	193	6%	4%	
Proximity pupils who attend other RC schools	41	27%	5%	17%
3km radius pupils who attend other RC schools	94	22%	15%	16%
School C = Outer London RC Girls School				
Current intake	150	3%	4%	
Proximity pupils who attend other RC schools	4	0%	25%	0%
3km radius pupils who attend other RC schools	47	26%	23%	0%
School D = Outer London RC Girls School				
Current intake	87	6%	12%	
Proximity pupils who attend other RC schools	14	14%	11%	7%
3km radius pupils who attend other RC schools	123	16%	25%	27%
School E = Outer London RC Girls School				
Current intake	175	1%	9%	
Proximity pupils who attend other RC schools	3	0%	50%	0%
3km radius pupils who attend other RC schools	22	0%	19%	36%

The above analysis is repeated in Table 5.11 for the five Church of England denomination schools that scored lowest on the FSM *proximity ratio*. Once again, for the most part, pupils living close to these élite CofE schools, but attending alternative CofE schools, are more likely to be FSM eligible or low ability than pupils in the élite school themselves. There are fewer CofE secondary schools than there are RC schools in London, so generally there are fewer pupils who live close to a CofE élite school but attend an alternative CofE school.

School F has an exceptionally small number of pupils who obtained low scores in their KS2 tests. By contrast, pupils who live close to school F but attend other CofE schools have FSM and ability characteristics more typical of Inner London. None of these pupils attend schools with superior peer groups.

School G has very low FSM eligibility rates and few low ability pupils. Pupils who live in the area but attend alternative CofE schools have high levels of FSM eligibility and are of a lower ability, on average.

School H has relatively few local pupils attending alternative CofE schools. The schools itself has few FSM eligible children and a slightly higher ability intake overall than the non-attending CofE pupils in the area.

Schools I and J are located some distance from alternative CofE schools, so few pupils living in the direct locality attend other CofE schools.

All five schools required the demonstration of religious adherence. In addition, three interviewed pupils and/or parents; two used banding (in one case this was banding skewed towards higher ability pupils); one selected a proportion of pupils on the basis of aptitude in languages; two gave priority to children of former pupils; one used 'compassionate' factors; and one gave priority to pupils with pastoral reasons. Only one school used none of these admissions criteria or practices.

5.7 Discussion

The analysis presented in this chapter shows that own-admissions schools have intakes that are more advantaged than community schools, even when the characteristics of local neighbourhoods are taken into account. Schools achieving intakes that are far more advantaged than their local neighbourhood almost always use explicit potentially selective admissions criteria as part of their admissions process. The more of these selective criteria are used, the more advantaged the school's intake appears to be. However, it is not possible to exclude parental choice as also playing a role in the production of these advantaged intakes.

Table 5.11: Cream-skimming analysis of 5 élite CofE secondary schools

	N	FSM	Low ability	High ability peers?
School F = Inner London CofE Girls School				
Current intake	90	12%	2%	
Proximity pupils who attend other CofE schools	4	25%	25%	0%
3km radius pupils who attend other CofE schools	44	23%	22%	0%
School G = Inner London CofE Girls School				
Current intake	151	5%	6%	
Proximity pupils who attend other CofE schools	24	33%	38%	13%
3km radius pupils who attend other CofE schools	157	40%	29%	12%
School H = Outer London CofE Mixed School				
Current intake	105	4%	15%	
Proximity pupils who attend other CofE schools	4	25%	25%	50%
3km radius pupils who attend other CofE schools	56	5%	11%	45%
School I = Outer London CofE Mixed School				
Current intake	197	9%	8%	
Proximity pupils who attend other CofE schools	4	25%	25%	25%
3km radius pupils who attend other CofE schools	51	22%	27%	47%
School J = Outer London CofE Mixed School				
Current intake	181	6%	17%	
Proximity pupils who attend other CofE schools	2	50%	0%	0%
3km radius pupils who attend other CofE schools	21	10%	11%	5%

The fact that explicit admissions policies have been shown to be an important source of school stratification suggests that the new School Admissions Code (Department for Education and Skills, 2007b), combined with better adherence to the Code by schools, should help to balance intakes across the schooling system. However, although explicit admissions criteria appear to be an important source of stratification, this does not exclude the possibility that more covert processes – such as marketing to affluent neighbourhoods and the use of complex supplementary forms – are not taking place within schools to help achieve more advantaged intakes.

Religious schools have been shown to be far more advantaged in their intakes than either Foundation or Community schools. The data suggest that if we take a Community school and a Religious school located in neighbourhoods with the same demographics, the Religious school might have as many as 50 per cent more top ability pupils than the Community school does. It is straightforward to attribute these advantaged intakes to the admissions policies of Religious schools, and in particular the processes used to establish religious adherence. Given that the religious background and practices of the family are a social phenomenon, it is not surprising that the religious adherence criteria is correlated with social class. The current measurement of religious adherence on a 'continuum' justifies the collection of family background data, giving religious schools the means to socially select pupils, should they wish to do so. It would seem inherent that problems will always arise in this process, because some questions deemed as relevant to establish religious adherence – such as marital status and place of child's baptism – by their nature reveal information about the social background of the family.

However, there is no proof that this covert cream-skimming is actually taking place in Religious schools: it is possible that more affluent families are able to score higher on the scale of 'religious adherence', either because they are more religious, or because they are better at meeting the school criteria for religiosity. Thus, it is possible that Religious schools are cream-skimming (that is, using admissions criteria that systematically favour more affluent families), but that they are doing so entirely inadvertently in the process of selecting by religious adherence.

An interesting policy question is whether religious admissions criteria need always favour higher income and better educated families. This would largely depend on the association between church attendance and individual characteristics such as education and income levels. Survey evidence from the British Household Panel Survey and the National Child Development Study suggests that better educated individuals are indeed more likely to attend church (Brown and Taylor, 2007; Sawkins et al., 1997). The evidence on the association between income and church attendance is less clear, with Iannaccone (1997) suggesting there is little relationship and Sawkins et al. (1997)

estimating that middle income earners are more likely to attend than low and high income earners. None of these measured associations are as pronounced as the observed differences in the propensities of children from different backgrounds to attend religious schools, which suggests that there are church-goers of lower social classes whose children do not currently attend Religious secondary schools. It is possible that they do not wish their children to attend, but alternatively they may not be negotiating the complex systems of school admissions successfully (because they lack skills to do so or do not value education highly enough to invest sufficient time in the process). For example, they may not attend the correct feeder primary school or a preferred local church for the correct number of sessions over the length of time specified by the school.

One way to simplify the admissions process for these families would be to establish a nationally agreed binary criteria of 'religious adherence' that families are deemed to have either met, or not met. Once this is established, Religious schools could then rely solely on the presence of a signature on a form from a religious leader to decide who has priority in the admissions process. This would avoid the need for the schools themselves to collect family background information, although it is still open to manipulation by parents. There is no easy solution to the problem of Religious school admissions, but a policy such as this could simultaneously remove the means by which covert cream-skimming is possible, while simplifying the admissions process for parents.

Chapter 6

The effect of school competition on pupil performance – a review of literature

This chapter sets out the theoretical framework that is used to argue that competition between schools can improve standards of teaching and pupil performance. It reviews the existing quantitative evidence on the effects of school competition on pupil performance, assessing the applicability of the international evidence to the English institutional environment.

In the context of schools, competition refers to rivalry between two or more schools for pupils. This competition may be over both number and type of pupils. With reference to the latter, schools may prefer to take pupils who have performed well at primary school or who come from higher social class backgrounds, using these as indicators of a pupil's likely future success in their education.

In England, the relationship between schooling provision and the neo-classical competitive market is weak, not least because there is no direct price for state education and schools do not profit-maximise. It is also true that peer group characteristics enter the education production function leading to the phenomenon whereby the characteristics of the demand for schooling affect the quality of supply. Furthermore, competition in local schooling markets tends to be competition amongst a few providers, making conventional analysis of markets with a large number of providers inappropriate (Davies et al., 2002). Despite these differences, the arguments for school reform over the last twenty years have been based on the neo-classical economic model of school competition, so it is useful to consider the model in detail, as it applies to the provision of schooling.

The first part of the chapter sets out the theory to explain a school's likely conduct in a competitive market. It also shows that the equations usually used to estimate the effects of school competition on aggregate performance in an area mask the separate effects of competition upon a school's teaching and organisational effort, changes in sizes of schools with differing effort levels and changes in the peer groups experienced by pupils due to sorting.

The final part of the chapter describes the findings of the English and international evidence on school competition, most of which are concerned with measuring the aggregate area effect of competition, as opposed to the distribution of individual or school-level outcomes.

6.1 Theoretical models of school competition

6.1.1 The school utility function

Schools are complex organisations, perhaps best described as a series of principal-agent relationships between parents, teachers, headteachers and governors or local authorities. A school's response to some competitive pressure will depend on the utility function of headteachers, teachers and governors of a school. We generally assume their collective utility at time t, U_{jt}, is a function of the size of school, n_{jt}, the growth or decline in pupil numbers between $t-1$ and t, Δn_{jt}, a measure of the pupil characteristics, x_{jt}, and the effort they must exert at work, e_{jt}:

$$U_{jt} = u(n_{jt}, \Delta n_{jt}, x_{jt}, e_{jt}) \tag{6.1}$$

The rationale for these variables entering the school's utility function are discussed below.

Teaching effort

The utility function we construct for the headteacher and teachers depends partly on our belief about whether these workers are principally self-interested or altruistic workers. Le Grand (2003, page 27) labels these respectively as knaves who are '*self-interested individuals who are motivated to help others only if by doing so they will serve their private interests*', and knights who are '*individuals who are motivated to help others for no private reward, and indeed who may undertake such activities to the detriment of their own private interests*'. Most teachers are not pure knights or knaves, but operate somewhere on a selfish-atruistic continuum. The attitude of any

one particular teacher can only be treated as partially fixed since they will likely have a motivational response to their environment and interactions with other teachers.

If teachers are primarily knaves, and if they derive a negative utility from putting effort into their work, they will engage in activities that minimise effort and make their working lives more comfortable. This might include altering their teaching style and minimising the amount of marking of class and homework they do. Collectively they would support school policies that seek to maximise rents for the employees (for example, via providing comfortable staff facilities and small teaching loads). However, it is perfectly possible that teachers are knaves, yet derive pleasure from certain teaching activities, and so minimise some, but not all, effort in their work. They may, for example, work hard in the classroom, but not in administrative aspects of their work.

Alternatively, teachers might be altruistic to the extent that they care deeply about pupil outcomes, regardless of the incentive structure that exists. However, although these knights might be diligent and work hard, they may not focus on the things that are important to parents.

Number of pupils and school growth

Assuming a pupil-led funding mechanism (as was introduced in the UK under the Education Reform Act 1998), growth of pupil numbers is likely to be the primary goal of schools since it equates to more funding. This assumes that heads and teachers derive some utility from this increased funding, for example through higher salaries, superior resources/facilities, or the prestige of running a large school. This incentive is, of course, also contingent on the school having spare capacity in existing buildings or access to capital to increase capacity.

It is the inter-relationship between pupil growth and its effect on pupil background characteristics that explains whether school growth is likely to be a goal of schools, especially where a school is given no control over which additional pupils it admits as it grows. Epple and Romano (2003) use a theoretical model to argue that under a school choice program where places for out-of-neighbourhood pupils are allocated by lottery, few schools will participate in the programme. They do so by making the plausible assumption that pupils will wish to transfer from a neighbourhood school with low average peer ability to another school with higher average peer ability. The prospective recipient school will not know the actual ability of the non-neighbourhood pupil, but can rationally assume that, on average, they will be of lower peer ability than the incumbent student on the basis that they live in a lower average ability neighbourhood. Therefore, the clientele of the prospective recipient school should resist

accepting this student who would, on average, lower the peer ability in the classroom. Epple and Romano note that in practice in the US, in states where participation in choice programs is voluntary, high-income districts typically opt out, or 'ensure' they have no spare capacity. The fact that only 120 schools applied to increase their capacity in the UK between 1999 and 2005 is consistent with Epple and Romano's model (Select Committee for Education & Skills, 2005).

Whilst it is not clear that there is always an incentive for a school to grow, there are much clearer reasons why a school needs to avoid a fall in pupil numbers, leading to a utility function with respect to pupil numbers as shown in Figure 6.1. Given pupil-led funding, falling school rolls almost always lead to teacher redundancies and severe budget cuts in departmental spending, with an associated loss of morale for remaining teachers. In the extreme, a falling pupil roll would lead to school closure with loss of jobs for all, including the headteacher. However, it is not clear how serious this closure threat is across different countries. In Sweden, for example, schools with declining numbers have been forced to dismiss teachers and make budget cuts, but do not generally close (Skolverket, 1996, 1999). In the UK, just 90 schools (both primary and secondary) closed between 1995 and 1999 (Gorard et al., 2002a). In response, the government has attempted to increase this threat via the 'Fresh Start' programme, which aimed to close failing schools and open new ones (that now must be academies) on the same site. However, in these circumstances most teachers will retain their jobs (with the installment of a new headteacher).

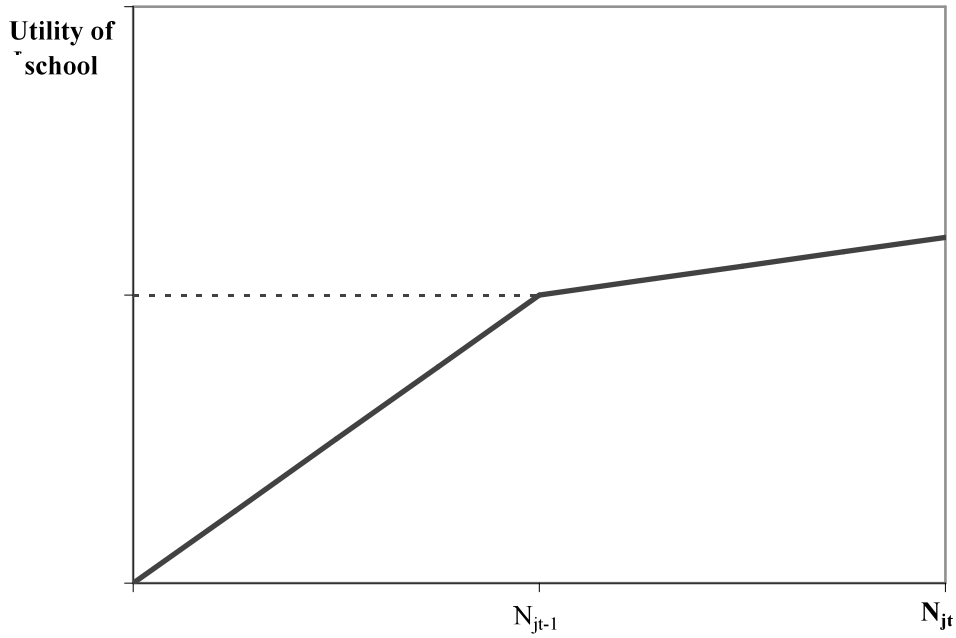

Figure 6.1: Utility function of school with respect to pupil numbers

Quality of pupils

The quality of the pupils at the school, meaning the distribution of prior ability and socio-economic background of the pupils, might enter the school's utility function for three reasons. First, higher SES and ability pupils might be easier or more pleasurable to teach and from the headteacher's perspective this makes the recruitment and retention of teachers easier. However, as previously discussed, teachers' motivations may be mixed and some may be attracted to working with children from a range of backgrounds. Second, performance appraisal remains contingent on actual pupil performance, despite the introduction of value-added measures of school performance which adjust for pupil intake. Lupton (2005) reports the strong association between Ofsted assessment of school quality and the social background of the pupils. Also, in June 2008, the government announced a 'National Challenge' to intervene in schools where fewer than 30 per cent of pupils achieved 5+ A*-C at GCSE, regardless of school value-added (Department for Children, Schools and Families, 2008b). Third, in the long-run, a high ability pupil intake ensures the school's survival, since they in turn attract other pupils to the school, assuming that parents value the league table position of school more than value-added as a measure of quality (as the evidence in Chapter 3 suggests).

6.1.2 The structure of competitive markets for schools

The next two sections distinguish between the classical concept of structural competition and competition as conduct and behaviour on the part of schools. This approach follows the structure-conduct-performance (S-C-P) model (Scherer and Ross, 1989), in which performance depends upon the degree and form of competitive conduct on the part of producers, in our case, schools. Competitive behaviour on the part of schools is, in turn, related to the structure which they face.

Although it is useful to distinguish between these aspects of competition, it is not the case that the structure of the schooling market is fixed. The structural features of a schooling market described here are likely to change in response to the conduct and performance of the system itself.

In describing the structural level of competition a school faces, it is useful to use the concept *residual elasticity of demand*, which '*measures the change in demand for school places in response to a change in [only that] school's quality*' (Bayer and McMillan, 2005, page 8). The determinants of a school's residual elasticity can be broken down into the availability of close substitutes in the marketplace (the size of the parental choice set and the diversity of schools within the choice set) and the responsiveness of households who are served by the school to changes in the school's performance. A

school's performance should generally be taken to mean a school's league table position, given the evidence on parental choice presented in Chapter 3.

Size of the choice set

The amount of competition in an area will depend on the size of parents' feasible choice set. The size of the parental choice set is a function of both the household attributes and location attributes, as discussed in Chapter 3. Determining the size of a choice set for any particular family is very difficult and subjective. We can define a feasible choice set given a fixed location as being *the set of schools that it is feasible for the child to travel to and that would be prepared to admit the child, given the household's location and attributes and holding constant the preferences of all other households.* So, the size of the choice set will be a function of population density, the costs and availability of transport, and the admissions policies in operation and ease of collecting information on schools. This type of definition emphasises that the households with the largest choice sets will have a higher ability child (who can pass entrance tests to selective schools), a religious affiliation and will live in an urban area. However, this doesn't take account of the capacity of households to choose schools by moving house. Indeed, in very rural areas, in areas without good public transport and where strict catchment areas are in place, this is the only means by which parents can choose. However, choosing by moving house is much more difficult than other forms of choice, and fewer families are willing to consider it, so it will be the case that competition is higher where the feasible choice set given a fixed location is high.

Diversity of schools within the choice set

The diversity of schools can be described in terms of hierarchical diversity, meaning large differences in the league table position of neighbouring schools, whether due to peer groups or genuine differences in effectiveness, and non-hierarchical diversity, meaning characteristics that distinguish the school from its competitors such as specialist sports, music facilities or religious ethos. In general, the greater the level of diversity of schools, the lower the residual elasticity of demand and thus structural competition. For hierarchical diversity, if a school changes its effort level and increases pupil performance, this will attract few extra pupils to the school if this extra effort does not close a very large gap between the test scores of it and the next best school in the league table. For non-hierarchical diversity, parents may make choice decisions based on idiosyncratic preferences for school features that are unrelated to test score; therefore effort to improve test scores may still not attract them to the school.

Taylor (2001), in a study of six English LEAs, argues that relatively few schools operate with the presence of non-hierarchical competition. This would be consistent with the general equilibrium models presented in Chapter 3 that imply there is instability in non-separating equilibria where schools have the same peer group quality. Of course, if schools start from a position of similarities in exam performance, for a period of time competition should be effective as school popularity will be very sensitive to small changes in performance relative to peers. This was observed by Fiske and Ladd (2003) in New Zealand where competition worked best for schools who started on a 'level playing field'. However, once a school improves its exam performance above those of neighbouring schools, the general equilibrium models of school choice suggest it will attract the higher income families from other neighbourhoods, thus improving its peer group, which will significantly raise its league table position further above its rivals. Once this has happened, the competitive threat the school experiences will fall.

Responsiveness of households to changes in school quality

A third determinant of a school's residual elasticity of demand is the responsiveness of households in the area to a change in a school's performance. For a strict neighbourhood schooling system, the key parameter is the price elasticity of demand for school quality because this tells us how much households value schools in their utility function (see Chapter 3). There may be heterogeneity in this price elasticity, so school located in low income areas may have lower structural competition, *ceteris paribus*, than a school located in a high income area where parents are very sensitive to changes in school quality. Also, where households can exercise choice given a fixed location, a key parameter is a measure of the school quality-distance to school trade-off, i.e. an elasticity of demand for school quality with respect to journey time.

The evidence presented in Chapter 3 suggests that 'school quality' can mean different things to families of different social classes, but that most parents proxy 'quality' as the school's 'raw' league table position without taking proper account of the progress their child is likely to make in the school. This assessment of the combination of school effort and peer characteristics means that parents are less responsive to changes in school effort than they would be if they were able to measure true effort of the school directly. Also, their use of league tables to assess school quality means that they are observing a lagged measure of quality based on the educational experiences of children 6 years older than their own, which bears little correlation with the expected school quality that their child will experience (Goldstein and Leckie, 2008). This again lessens the incentive for schools to focus on effort.

Policies concerning school capacity and closure

A final structural feature of a school market that affects competition concerns how real the threat of closure is. Where a local school-aged population is growing faster than capacity growth, the closure threat may be very low indeed, regardless of the performance of the school. This has been the case in many places where market reforms were introduced (e.g. UK between 1988 and 2005; California from late 1990s to present). Three policies can increase structural competition via the manipulation of capacity. First, government can intervene and close schools down, even if they still contain pupils. This type of policy is certainly not a market reform, but might be deemed necessary in areas where the responsiveness of some parents to school quality is very low. Second, government can provide existing schools with capital to increase capacity and can subsidise the entry of new schools to the market. Third, governments can force popular schools to increase their capacity, which is likely to be necessary given models predict that capacity increases at the most popular schools are, on the margin, likely to lead to a deterioration in peer quality.

The location of spare capacity in the system is crucial to a consideration of how effective it is in producing a competitive threat to schools. The problem with introducing new capacity into a system, regardless of where it is initially located, is that it very quickly clusters in the school at the bottom of the local market hierarchy and will then never again be used (unless the population increases). So, this spare capacity becomes redundant and presents no further threat to other schools in the area.

6.1.3 The conduct of schools in a competitive market

Internal efficiency and test scores

One response to a competitive market is to improve the internal efficiency of the school by focusing effort on increasing the school 'output' given its fixed resources (and holding the quality of pupils as inputs constant). They might do this simply through increased effort, or by using resources more innovatively or efficiently. Headteachers may introduce monitoring and benchmarking systems for teachers to induce internal competition. They may also hire teachers more carefully and be more prepared to fire teachers who are underperforming. Finally, they may simply encourage an ethos of academic performance at the expense of other outcomes. This would be a problem if the educational outputs valued by children and society were much broader than exam performance and included, for example, consumption benefits (the enjoyment of the education process itself); investment benefits (higher pay, job satisfaction and enjoyment of leisure); and various external benefits, including shared values.

It is difficult for entities as complex as schools to demonstrate accurately their quality to potential customers and creating indicators of school quality is difficult (Adnett and Davies, 1999). In England, league tables have published the proportion of pupils attaining 5+ A*-C at GCSE at 16 for each school since 1992. This led to an increased focus by parents and schools on this metric of success, with some unintended consequences. First, schools increased resources directed at the sub-set of pupils on the C/D borderline, potentially diverting effort away from pupils at the top and bottom of the ability distribution (Burgess et al., 2005a; Wilson et al., 2006). Second, schools focussed on subjects where the return to increased effort was greatest. Some schools used GNVQs worth 4 GCSEs to boost their pass rate and others diverted attention away from maths and English, which are the most important qualifications from the perspective of the child, but were weighted equally with other subjects in school league tables (this problem is now resolved because for league tables from 2006 onwards, the 5 A*-Cs must include maths and English).

Other evidence suggests teachers engage in activities likely to be noticed. For example, Woods et al. (1998) find that school 'strategies' tend to emphasise academic outcomes and physical facilities which are highly visable, even though parents put just as much emphasis on, for example, discipline and school climate. An US empirical investigation of the effect of intra-district choice on teacher incentives suggests that choice does enhance teacher motivation, though when teachers become more concerned with how their efforts are perceived in the marketplace they tend to concentrate on activities that are more likely to be noticed by parents, and not on tasks like classroom preparation (Rapp, 2000).

Cream-skimming

Cream-skimming occurs where an oversubscribed school selects pupils of high intrinsic ability so as to improve the school's performance in examinations at relatively low cost (Le Grand, 2003, page 10). They have a clear incentive to do this where parents focus on a metric of raw league table performance, though if these students are also easier or more pleasurable to teach, this incentive also exists to some extent without a market for schooling (Thrupp, 1999).

There are two methods by which schools can cream-skim. The first is known as overt cream-skimming and it takes place through the use of admissions policies (or oversubscription criteria) to systematically target and accept student of higher ability or higher SES into their school. For example, in the UK in the 1980s there is some evidence that the admissions structure of City Technology Colleges (CTCs) allowed 'cream skimming'. These schools were required to take a sample of children whose ability was roughly representative of those applying. However, only highly motivated

parents applied: they had to seek out the application form, arrange for their child to take a nonverbal reasoning test, attend an interview and confirm that their child intended to continue in education to 18. Some claimed that CTCs were systematically selecting pupils whose parents appeared keen and enthusiastic at interview (Walford and Miller, 1991).

The West et al. (2004) analysis of comprehensive secondary school admissions criteria in England reveals a significant minority of schools using criteria which appear to be designed to select a certain group of pupils and so exclude others. Despite all the schools in their study being 'comprehensive' schools, West et al. found the following opportunities to select:

- Admissions criteria relating to **ability/aptitude** were mentioned by 11.2% of Foundation schools, 6.5% of Voluntary-Aided schools and 0.3% of Community schools.
- School-administered ability **banding** by schools was found to be more prevalent in Voluntary-Aided/Foundation schools (5%) than Community schools (2%).
- Giving priority to children of **employees/former pupils** etc. is now against the school admissions Code of Practice (Department for Education and Skills, 2003), but 20% of Voluntary-Aided/Foundation schools (versus 5% of community schools) continued to use this criterion in 2002.
- 13% of schools in West et al.'s sample used **religious** criteria to admit pupils.

A second form of cream-skimming comes via more 'covert' selection and Gewirtz et al. (1995), in a study of London schooling, argued that subtle 'social targeting' is more prevalent than overt cream-skimming in England. Social targeting means:

> giving priority to promoting the school to middle-class and more able pupils (often done through geographical targeting, focusing on feeder schools located where there are concentrations of middle class and more able pupils). Such strategies also include: attention to the 'traditional trappings' of the school attention to the schools achievement of A to C grades at GCSE; specialisation in a 'strong' curriculum area; downplaying the level of rhetoric at least of the caring, pastoral side of schooling; and generally sending signals indicating what sort of parents are made most welcome.
>
> (Woods et al., 1998, page 159)

There is widespread evidence of an emphasis upon promotional image and image enhancing changes following 1988 in the UK (Gewirtz et al., 1995). Some of this social targeting constitutes clear rivalrous behaviour towards other secondary schools, for example, by marketing to primary schools in their catchment area and by running

school coaches from their catchment area neighbourhoods (all reported by Levačić, 2004). Thus, if we believe that schools face a clear incentive to behave in this way, it is likely to reduce the level of collaboration and innovation between neighbouring secondary schools (Levačić reports that collaboration was more common between distant secondary schools that were not in competition with each other).

Diversification

One means of reducing the level of structural competition a school faces, is to diversify its product and therefore reduce its residual elasticity of demand. The decision by a school to adopt this approach might be taken alone, or it might be part of a collusive strategy with other schools to segment the local marketplace, thereby reducing competitive pressure. Ladd and Fiske (2003) claim this was a common response to competition in New Zealand, where schools re-branded themselves as specialist arts or sports colleges, for example. However, it is not always the optimal strategy to adopt this approach if parental tastes for schooling are essentially quite homogeneous, i.e. they value league table position above all other attributes of the school. In the UK, despite the introduction of a Specialist Schools policy (where secondary schools are given extra funding to 'specialise' in a particular subject), evidence shows that school choice reforms appear to have reinforced local schooling hierarchies, reduced the diversity of provision and increased differences in the mean level of academic attainment between schools (Adnett and Davies, 2002).

This seeming homogenisation of schools that is witnessed as schooling markets become more customer-focused and market-orientated is interesting given that an alternative competitive response would be to adopt a niche strategy, protecting the school from competition (Pardey, 1991; Saltman and Von Otter, 1992). It is possible, but unlikely, that the homogenisation can be entirely explained by the greater regulation of schools, particular with respect to curriculum content (e.g. the KS3 strategy). Alternatively, the reduction in the diversity of provision may be a competitive response to either (1) parental preferences that are essentially homogeneous; or (2) a 'Hotelling'[1] style response to non-homogenous parental preferences (Hotelling, 1929). The weight of the current evidence on parental preferences, presented in Chapter 3, would seem to weight towards the former. In other words, greater choice in a market does not necessarily produce greater variety, which is one of the purposes of introducing choice in the first place.

[1]This refers to analysis of the effect of absolute and relative positioning of firms (or schools) on their success in a market. It demonstrates the circumstances under which the optimum strategy for all firms (schools) would be to meet the preferences of the median buyer (parent).

6.1.4 The dynamics of a competitive marketplace

Well-functioning competition in the quasi-market cannot exist if schools in certain circumstances find it very difficult to improve their position in the local hierarchy of schools, regardless of effort or quality of education offered. In a longitudinal study of six LEAs between 1991 and 1998, Levačić and Woods (2002) found that schools with lower concentrations of social disadvantage relative to other local schools were able to improve their GCSE results the fastest. Furthermore, schools with high concentrations of social disadvantage were liable to suffer a dual handicap as their relative social disadvantage tended to worsen over time. There is some evidence to the contrary. Glennerster (2002) points out a narrowing of the gap in English KS3 scores between the top and bottom quartile of schools from 32 points in 1995 to 26 points in 2000. On this measure there was little change in the gap in maths and science, but certainly no widening gap. Furthermore, the schools with over 40% of pupils on free school meals have been catching up with the schools where less than 5% have free meals. One caveat about these figures is important: the schools at the top are already near the achievement targets and now have little room for improvement, thus 'catch-up' could be an artefact of truncating the measurement of improvements with fixed targets.

6.2 School competition – the estimation problem

Empirical studies of the effects of school competition on mean pupil performance usually employ an aggregate area-level measurement of the 'effect' of competition, which captures the total effect of three mechanisms:

1. increases in school efficiency at one or more schools, induced by competitive pressure;
2. changes in the relative sizes of schools, with more effective schools possibly gaining additional pupils; and
3. student sorting effects, which alter the distribution of pupils across schools, which might either induce better matching effects or beneficial changes in peer group effects.

Formally, using an example of two equal-sized schools, A and B, operating in a town that are initially not in competition (because, for example, the authorities control admissions and capacity at the schools), a child's educational outcomes, y, at $t = 0$ are:

$$y_{i,A}^{t=0} = \beta_0 + \beta_1 \theta_A^{t=0} + \beta_2 C_A^{t=0} + \beta_3 U_A^{t=0} + \beta_4 x_A^{t=0} + \epsilon_A^{t=0} \tag{6.2}$$

$$y_{i,B}^{t=0} = \beta_0 + \beta_1 \theta_B^{t=0} + \beta_2 C_B^{t=0} + \beta_3 U_B^{t=0} + \beta_4 x_B^{t=0} + \epsilon_B^{t=0}$$

In other words, we assume that a child's test scores is determined by some function of the peer group it experiences in the school, θ; the effort the school puts in as a direct result of competition, C (where we can assume that $C^{t=0} = 0$); the school effect on outcomes not attributable to peer group or competition (e.g. because of strong leadership), U; and the known, x, and unknown, ϵ, characteristics of the child.

Then suppose in time period $t = 1$, open enrolment is introduced and schools may compete freely for pupils. As a result, the $2n$ pupils in the town are redistributed across schools and A grows to size $n_A > n > n_B$ (retaining $n_A + n_B = 2n$). The effect of this competition on mean pupil performance in the town ($\sum_{i=1}^{2n} y_i^{t=1}/2n - \sum_{i=1}^{2n} y_i^{t=0}/2n$) is the aggregate of three effects:

1. $\beta_2 C^{t=1} > 0$, assuming that the impact of competition from school A on school B is identical to the impact of competition on school B from school A;[2]

2. $\frac{\beta_3}{2n}[U_A(n_A - n) + U_B(n_B - n)] > 0$, but only if school A was more effective than school B in the first period and continues to be so, so that parents are correctly identifying and choosing the more effective school (given its inputs). There is no reason to expect this to be the case where 'raw' league table results are published.

3. $\frac{\beta_1}{2n}[n_A\theta_A^{t=1} - n\theta_A^{t=0} + n_B\theta_B^{t=1} - n\theta_B^{t=0}]$, which may be positive, negative or zero. If we take the common assumption that the peer group $\theta = \frac{1}{n}\sum_{i=1}^{n} x_i$ (the mean characteristics of the school), the marginal gain by any one student represents an equal marginal loss in peer group by another student. Therefore this change in the peer group distribution has no effect on mean pupil performance. This assumption may, or may not, be reasonable: the Bayer and McMillan (2005) findings (described in this chapter) imply that there are decreasing returns to a good peer group (so segregation reduces mean pupil performance). By contrast, the UK literature on ability streaming suggests that the benefits to high ability pupils outweigh the disadvantages to the lower ability pupils (e.g. Galindo-Rueda and Vignoles, 2004; Atkinson et al., 2006).[3]

[2]We generally assume that the effect of competition cannot be negative, although the theoretical models of De Fraja and Landeras (2006) suggest that the strategic interaction among the participants in the education process may make incentives backfire and competition have perverse effects. For example, students may reduce their effort when teachers increase theirs, and school attempts to attract the better students may plausibly reduce students' attainment

[3]In any case, our ability to make judgements that compare the size of the gains for high ability pupils with the size of the loss for lower ability pupils depends on the construction of a scale of outcomes. For example, is an improvement from grade B to A equal to the improvement from grade F to E? So, the overall direction of the effect on mean test scores is solely a function of the construction of this outcome measure, where the construction of cardinal intervals is necessarily quite arbitrary.

6.2.1 Key measures of competition between schools

The key explanatory variable of interest – the amount of competition – often proxies some aspect of structural competition in the market. This is because structural competition tends to be more amenable to measurement and more readily subject to manipulation by policy-makers. The first set of variables used by empirical studies of school competition measure structural competition at the area level in terms of:

- Number of schools in area;
- Share of a particular school type enrollment in area, e.g. private school, religious school or Charter school;
- Herfindahl index of school competition in area;
- Population density.

There are two problems with the construction and use of these measures. First, the 'area' must first be defined. If it is too large an area, the competition measure will be a poor proxy for the competition actually experienced by schools. If it is too small an area, it will fail to capture the full extent of competition between certain schools. The problem tends to be that some schools exist within very small competition areas, but others (such as religious schools) will have large recruitment areas. In practice, typical areas used include local authorities, municipalities and non-overlapping areas of competition defined by the researcher for the purpose of the study. The second problem with these measures is that, except for population density, they are all endogenous to the market under investigation. For example, the share of pupils in private schools will be partially determined by public school quality. The studies described often overcome this endogeneity by instrumenting the competition measure with an exogenous variable. Population density would seem to be a strong competition measure since it is exogenous. However, there are alternative explanations as to why urban areas may systematically differ, for example schools may be more likely to collaborate and they have access to a different type of teacher labour market.

The second set of measures of competition are also endogenous structural measures, but they identify the extent of competition at the level of the individual school (although they can also be aggregated up to area level):

- Distance to school to nearest 'choice' school (e.g. Charter school or private school);
- Number of other schools within a 1, 3 or 5 km radius;
- Number of other schools within a 10 minute drive-time zone;
- Presence of a 'choice' school within a 5 km radius.

The third set of measures are also school-level structural variables, but these exploit the availability of data regarding pupil geo-locations, for example:

- The proportion of pupils living within 5 km of the school that have switched to a 'choice' school;
- The mean level of choice a school's pupils have, measured as schools the pupil can reach within some defined time or distance.

Fourth, an endogenous measure is used that attempts to capture the competitive threat as measured by (lagged) school exam performance of other schools in the area. Finally, one study uses survey data to measure the headteachers' perception of the level of competition in the area.

6.3 Literature on competition and pupil performance

This section describes and evaluates the key empirical studies into the effect of competition on pupil achievement from across the world. All the main UK studies are analysed, and in addition recent international studies are described, providing they use high quality datasets with clear methods for identifying the effect of competition.

6.3.1 English non-experimental evidence

The simultaneous country-wide implementation of legislation arising from the 1988 Education Reform Act, combined with a lack of historic pupil-level national datasets, makes evaluating the effect of these reforms very difficult. The non-experimental studies all rely on exploiting structural differences in the level of competition between schools in different geographical locations. However, as discussed in the previous section, almost all 'structural' measures of competition are also endogenous.

Bradley et al. (2001) use Data Envelopment Analysis (DEA) to calculate the efficiency frontier for all secondary schools in the UK, where the school's inputs are the proportion of pupils ineligible for free school meals (FSM) and the proportion of teachers who are qualified at the school. Individual school efficiency, in terms of the proportion of pupils gaining 5+A*-C at GCSE and school attendance rates, is measured in relation to the efficiency frontier. School competition is measured as the number of other schools within a given kilometre radius (<1 km, 1-2 km, 3-5 km). In cross-section they find that schools facing competition from non-selective schools are more efficient. Between 1993 and 1998 they find that the change in relative efficiency is positively related to competition (particularly the number of schools within 1 km). In a related paper

they also show that change in test scores is positively correlated with test scores of other schools in the same district and interpret this as a competition effect, although common gentrification of areas would produce the same association (Bradley et al., 2000).

Bradley et al. report issues with the quality of the school-level data in this study: the validity of the attendance measure is in some doubt (in relation to reported truancy rates in the Youth Cohort Study); there are very few school-level controls and no pupil-level background variables or prior attainment. Furthermore, their competition measure is highly correlated with population density, so the study cannot distinguish competition effects from other effects associated with cities.

Ros Levačić (2004) also tries to identify an association between a structural measure of competition (the number of schools within a non-overlapping approximate area of competition) and school performance in Annual Schools Census data, but finds no association over the slightly longer period between 1991 and 1998. There are several possible explanations for this apparent contradiction with Bradley et al.'s work. First, she has a much smaller sample of just 220 schools in six LEAs. Second, her structural measure of competition is quite different and arguably weaker since it relies on non-overlapping competition spaces and is not a school-level measure. Third, she employs an OLS regression approach rather than a DEA. DEA is potentially superior to regression where organisations have clearly defined multiple outputs, and particularly where increasing one type of output is likely to reduce another type of output. However, in the case of schools it is less clear that attendance rates and exam performance should be treated as distinct outputs since improving attendance rates is likely to improve a school's exam performance. In addition the attendance variable data is poor, so using standard regression analysis for schools seems sensible.

The more interesting aspect of Levačić's study is the survey of 220 headteachers asking them how much competitive pressure they perceived their school to be under. She found no association between her structural measure of competition and perceived competition by headteachers. However, schools where the headteacher had a high perception of competition (reported more than 5 competitors) outperformed other schools with 4.0 - 5.5 percentage points more 5+A*-C grades and a 0.3 percentage point/per annum greater increase in this measure from 1991 to 1998. It is not clear that this finding can be interpreted as a causal relationship between competition and school performance: it is equally likely that headteachers who are intrinsically highly motivated by league tables and other outcome measures are more likely to perceive they are in a competitive environment.

In a recent study, Bradley and Taylor (2007) use school-level data over a very long time period (1992-2006) to analyse the effect of competition within local authorities (LAs)

on the proportion of pupils achieving 5+A*-C at GCSE. The competition variable is measured by the average exam performance of all schools in the LA, lagged one year, excluding the school in question. A Herfindahl index of concentration is also used as a measure of competition in the LA. The estimation strategy employs school fixed effects and time-variant school control variables to attempt to deal with pupil sorting, although the limited range of school-level pupil characteristic variables available for this time period means that sorting based on unmeasured characteristics might confound results. The study finds that a one percentage point increase in the exam performance of competing schools is associated with an improvement of 0.38 percentage points in metropolitan schools compared to only 0.12 percentage points in non-metropolitan schools. However, once again, the gentrification of English cities would explain this type of result.

These associative relationships between structural measures of competition and pupil performance have been analysed in the more recent pupil-level data by Gibbons et al. (2008b), who show that secondary school progress is faster in areas of relatively dense population, but that the elasticity of school performance with respect to population density is low at +0.02 (Gibbons et al., 2008b). There are many different specifications in the paper that try to deal with the limited number of pupil controls in NPD. For example, one specification does this by using a dependent variable of pupil progress between 11 and 16, relative to the mean progress of primary school peers. The parameter of interest is the change in urbanness of school experienced by the pupil as a result of the primary-secondary transfer, relative to other pupils the child was at primary school with. One underlying assumption for an unbiased estimate of this parameter is that the change in urban density between primary and secondary is not correlated with unobservable individual characteristics. An example of when this might be violated would be an observation that more prestigious secondary schools tend to be located in town centres (since they are the former grammar schools). They cannot completely reject sorting on observables in this study (notably schools in urban areas have accelerating pupil progress), so their findings should be only suggestive of a causal relationship. They rightly point out that it is impossible to distinguish between competition and collaboration effects in urban areas.

The final non-experimental study to look at competition effects analyses the relationship between KS1-KS2 progress in the latter part of primary school and a school competition index in London (Gibbons et al., 2008a). The construction of the school competition index is complex. The first stage is to calculate the distance travelled by the median pupil at each school. Next, they calculate the size of the choice set of each pupil by asking how many within-LA schools the pupil can reach if s/he travels no more than the median travel distance for each school and remains within his/her local authority (LA). This definition of 'choice' is not structural but is actually endoge-

nous to actual choices currently exercised by parents, given the constraints imposed by school capacity. For example, if a set of parents are currently choosing to travel a long distance to a non-neighbourhood school, this will increase the choice measure. The competition index is calculated for each school by computing the mean of this pupil choice index for all pupils at the school. Since the resulting competition index is also endogenously determined by demand for schooling, an instrument of distance of school from the LA boundary is used to predict it, on the basis that this is exogenous to school performance. The instrument is, by construct, strong because the measure of the choice set restricts choice close to boundaries. They find no overall effect of choice or competition on school performance, though some positive effect of competition from religious schools.

It is possible to interpret this finding as showing that primary schools do not respond to competition from other schools. However, there are several serious questions over the validity of their competition measure that need to be considered further. First, once again underlying it is a structural approach that assumes that schools feel competitive pressure from schools in close proximity, but this is not necessarily so if the market for schools is very fragmented with neighbouring schools operating in very different competition spaces. For example, a full capacity school with a high income peer group is unlikely to perceive it is in competition with a partially full school with a low income peer group. Second, the identification strategy relies on the assumption that pupils do not generally cross an LA boundary in London and so schools experience no competitive threat from schools in other LAs. It is not particularly difficult to attend a school in a neighbouring LA, but it is true that only 4.7% of pupils do so in Gibbons et al.'s dataset. It is not clear whether or not the figure of 4.7% is sufficiently low to claim that boundaries significantly impact choice. Even if we could show that some parents were discouraged from considering schools in other LAs, provided high-income, motivated parents saw it as no barrier to choice, schools would experience competitive pressure across LA boundaries since these are the type of families that they arguably compete for. The validity of their competition measure needs to be considered in the context of the following question, and it is not clear how it can be answered: 'if the boundary disappeared, would schools operating near the boundary experience an increase in competitive pressure from other schools'? Our knowledge of parental choice strategies suggest that the boundary is more likely to act as a constraint for working class families, which suggests that it is only schools who cater for lower income households that are likely to experience the increase in competitive pressure.

6.3.2 English quasi-experimental evidence

There are two UK studies that exploit policy reforms to implement a quasi-experimental approach to studying competition effects. The first is a paper that identifies the effect of a potential increase in competition resulting from schools pursuing autonomy under the Grant-Maintained schools legislation. This paper by Clark (2007) is described in some detail in the next chapter, so the competition finding is briefly summarised here. School governing bodies were required to hold a vote of parents of current pupils at the school before electing to become Grant-Maintained. Damon Clark (2007) exploits this vote and data on the proportion of parents voting yes by employing a regression discontinuity design to compare change in exam results for those schools where they 'just' won over 50% of the vote to those schools where they just lost with under 50% of the vote. Using data between 1992 and 2000 he finds that Grant-Maintained schools did not appear to exert a greater competitive pressure on their neighbours than schools who lost the vote and remained under LA control. It is not possible to use his dataset to disentangle whether neighbouring schools felt no threat from GM schools (this is possible given rising pupil rolls over this period, especially if neighbours felt the GM school had no plans to increase capacity) or whether his structural measure of competition (once again) does not reflect the true competitive environment in which schools operate.

Burgess and Slater (2006) also try to employ a quasi-experimental approach by using the administrative boundary change of 1988 in Berkshire to estimate the impact of changes to school competition on pupil progress between KS3 and KS4. Using a small sample of 390 schools, they estimate a difference-in-differences model of changes in school effectiveness on changes in competition between a cohort finishing secondary school in 1997 and another in 2002. The changes in competition result from the division of Berkshire into 6 LAs, with the presumption that schools close to the newly created boundaries see a greater fall in competition than those in the middle of a new LA. The paper benefits from the availability of pupil-level data (though the first panel has no pupil-controls other than prior attainment so a bias correction is used). It finds no effect of the fall in competition on pupil outcomes, but there are several methodological explanations for this null result. First, the sample is small with the variation in 'competition' being driven from very few schools. Second, they are unable to show any evidence that competition did actually change following this reform. Very high proportions of pupils cross an LA boundary to attend secondary school in this densely populated part of England. Indeed, there are even official catchment zones for schools that cross these boundaries. Each LA in Berkshire advertises clearly how to get sent a brochures for schools in other LAs and admissions applications are on a common form for all schools inside and outside the LA. Third, Berkshire's education system

was always particularly complicated because selective systems of schooling operate alongside non-selective systems and so it is not clear how much, for example, the grammar schools in Reading ever seriously felt competition from the comprehensives in Windsor and Maidenhead. Finally, even if mobility did fall, if the higher-income well-motivated parents are still choosing schools across LA boundaries then this should mean there is no change in competitive pressure for most schools since these are the very families they are competing for.

6.3.3 US public school competition evidence

The US empirical literature on school competition provides a more positive picture of an effect on school performance. The Belfield and Levin (2003) literature review of 35 wide-ranging studies found that between 36% and 67% of estimations of the effect of a cross-sectional measure of competition are statistically significant and positive (with less than 3% negative and significant). However, the effect sizes are quite modest and are not generally robust to changes in the specification of equations. The typical study found that a 1 s.d. increase in competition (measured on a Herfindahl index of school concentration) increases test scores by 0.1 s.d. and no effects on drop-out rates were generally found.

Most of the studies reviewed by Belfield and Levin are measuring associations without employing an identification strategy to identify a causal effect. The most controversial study that attempts to do this is Caroline Hoxby's (2000) article on competition (measured as the number of school districts within a metropolitan area) and pupil performance in 8^{th}, 10^{th} and 12^{th} grade reading and maths tests. Using 1994 data from the National Education Longitudinal Study (NELS) of over 6,000 students in 209 metropolitan areas she finds a positive, statistically significant effect of competition on pupil outcomes and productivity. She has a good set of pupil and district-level control variables and deals with the potential endogeneity of the Herfindahl index by instrumenting it using the number of large and small rivers in the metropolitan area. The rationale for this is that rivers form a natural geographical boundary such that the greater the number of rivers, the more school districts there are likely to be. As an example of the magnitude of her findings, in an area with maximum competition she estimates a 3.8 percentage point effect on 8th grade reading scores and 5.8 percentage point effect on 12th grade reading scores, compared to an area with no competition between school districts.

Hoxby's findings have been strongly questioned by Jesse Rothstein (2005) who documents several errors in Hoxby's data and code. He also demonstrates that the estimated choice effect is extremely sensitive to the way that 'larger streams' are coded. When Hoxby's hand count of larger streams is replaced with any of several alternative, easily

replicated measures, there is no significant difference between IV and OLS, each of which indicates a choice effect near zero. Rothstein appears to throw the strength of Hoxby's claim into some doubt: when econometricians employ an IV technique it is important that their estimates are robust to small technical 'tweaks', and hers appear not to be.

Using a completely different approach to the problem of identifying causal effects in non-experimental data, Bayer and McMillan (2005) use an exceptionally detailed dataset to estimate a complete model of demand for schooling in San Francisco. In San Francisco there are no choice programs, so families exercise 'choice' for schooling by moving neighbourhoods. This means that the elasticity of demand for school quality is revealed in house price data at district boundaries. Bayer and McMillan run a simulation that uses these elasticities to calculate a residual elasticity of demand for school quality at each school. This is defined as the responsiveness of demand (measured in house prices) to changes in own school quality. Using cross-sectional regressions of elasticity on school performance, they find a one s.d. increase in the competitiveness of the school's local environment within the Bay Area leads to a 0.15 s.d. increase in average test scores. Their finding relies on the validity of the initial school demand model specification, but does appear to be robust to various different specifications. Nevertheless, this approach does need to be validated via replication in another city, which is difficult given the need for very disaggregated data. They also find an association between the level of school segregation and school performance (segregation lowers competition between schools) and this result could be more simply replicated elsewhere.

6.3.4 US Charter school competition evidence

Charter schools are US government-funded schools that are not subject to many of the same rules and regulations as traditional public schools. They do not select by ability and have been set up with the intention of raising academic standards, in part by exerting a competitive pressure on existing public schools. The UK Academies programme was, to a large extent, directly modeled on US Charter schools.

The quantitative literature described here is slightly different from studies described earlier because it tends to estimate the effect of Charter school competition on pupil outcomes at traditional public schools (i.e. these are not whole-area effects). The problem with this approach is that some (but not all) studies are incapable of properly distinguishing between changes in public school effectiveness and changes in pupil sorting. So, given that Charter schools can often attract a slightly more affluent demographic away from traditional public schools, studies with poor pupil control variables are likely to estimate a negative effect of competition. A second problem

with many of these studies is that they estimate models as though the location of Charter schools is exogenous to traditional public school effectiveness. The direction of this endogeneity problem depends on the model estimated: there is likely to be a negative correlation between public school performance and the location of Charter schools; however, there may be a positive correlation in panel data between changes in public school performance and changes in Charter school penetration, via a simple 'mean reversion' of any school's effectiveness.

Estimates of the effect of Charter schools on traditional public schools from across the US have varied in both the direction and degree of the competition effect. Carr and Ritter (2007) suggest there are two likely explanations for this. First, Charter school laws differ across US states, and this has affected the very nature of the programme. For example, in some states Charter schools can be set up by independent bodies, whereas in others only the local school district is allowed to create Charter schools. Second, construction of the primary explanatory variable – the amount of competition from Charter schools a public school faces – has been inconsistent across studies. This has meant that, even where studies analyse Charter schools in the same state, estimates of competition effects have differed.

The state of Texas is the only area where all studies have identified a positive competition effect of Charter schools. Booker et al. (2005) find competitive pressure produces test score gains in a very large 8 year panel survey of high school pupils in Texas. Their measure of competition is the number of actual students in the district lost to Charter schools, which is endogenous to pre-existing school efficiency and therefore instrumented using a lagged competition measure (independence of the instrument is met, they argue, because they have a measure of prior achievement for every child). The study uses student fixed effects to control for unobservable pupil characteristics, and thus exploits movements of pupils between school districts (there is very high pupil mobility in Texas and they argue it is mostly for reasons unrelated to schooling).

Bohte (2004) also find a positive effect of Charter schools using Texan data from similar years to Booker et al. (2005) and a variety of measures of competition. They estimate that a one per cent increase in Charter school market share is associated with a 0.1 per cent increase in district passing rates.

North Carolina is one of two states where different research teams have drawn opposite conclusions about the effect of the Charter school programme. Using similar years of data for grades 3 to 8, Holmes et al. (2003) found a small positive effect of competition on public school test scores, whereas Bifulco and Ladd (2004) found no effect on maths achievement and a negative impact on reading where Charters were within 2.5 miles of a traditional public school.

Similarly, the studies from Michigan state also conflict. Hoxby (2001), using difference-

in-difference means tests between areas with over 6% Charter market share and those with less found 4th and 8th grade reading and maths scores increased by a modest amount (she also found this in Arizona). However, Bettinger (2005) finds no robust evidence that test scores in neighbouring public schools increase as the number of Charters in Michigan increases. Unlike Hoxby, he can deal with the endogeneity of the location of Charter schools by instrumenting it using proximity to a state university (since these could set up Charter schools). However, his sample is of Charter schools is quite small (33 out of 1,800 schools in total) and this might explain his null result. In a much later study, Ni (2007) is able to look at the longer run effects of the policy in Michigan. He finds a negative impact on student achievement and school efficiency in traditional public schools. The effect is small or negligible in the short run, but becomes more substantial in the long run, which is consistent with the conception of choice triggering a downward spiral in the most heavily impacted public schools.

Studies from Milwaukee, Florida, California and Ohio have estimated no effect, or a slightly negative effect, of Charter schools on traditional public schools. Sass (2006, Florida) and Buddin and Zimmer (2005, California) both use school fixed effect regressions to look at the association between changes in competition from Charter schools on changes in school performance. Sass finds no effect on reading and only a small effect on maths achievement. Buddin and Zimmer find no effect on either, with a survey of headteachers reporting that they perceived little or no impact on their school. However, during the implementation of the policy, California was experiencing an increasing school population, which might explain why the new capacity introduced via Charter schools failed to impact on traditional public schools.

Greene and Forster (2002) find no effect of Charter school competition on test scores in grades 4 and 8, but a small positive gain for 10th grade students where a Charter school is within 1km in Milwaukee, Wisconsin. However, the identification strategy is weak, relative to other studies, since it estimates a school-level OLS regression with only basic control variables. Carr and Ritter (2007) uses a pooled times series regression design (with lagged outcome and school context variables as controls) to find a small but consistently negative effect of Charter schools on proficiency passage rates of nearby traditional public schools in Ohio. However, the lack of pupil characteristic control variables means they cannot exclude the explanation that this is a pupil sorting effect, rather than direct negative competition effect.

6.3.5 International private school competition evidence

Public schools do not just experience competition as a result of other public schools; some public schools compete for pupils with private schools. However, there are reasons to think that this competition from private schools is unlikely to be as effective

as competition from other public schools. This might be because choice of private school has little to do with the performance of public schools. Alternatively, the small number of parents actively choosing between private and public schools might mean that competition is never high enough to estimate a significant effect.

Caroline Hoxby (1994) uses the same dataset as described in her study above to examine the relationship between the proportion of pupils in private schools in the metropolitan area with the performance of pupil school pupils. Because the number of private schools in an area is clearly endogenous to demand for schooling, an instrument is needed for this explanatory variable. Hoxby uses the population densities of nine major religious denominations as her IV. It predicts the level of private schooling in areas because public schools must be secular, so religion communities were the principal drivers of the growth of private schooling in America. Hoxby argues that the incidence of religious groups is largely a historical phenomenon that is exogenous to her dependent variable (8^{th}, 10^{th}, 12^{th} grade reading and maths scores), provided the religious denomination of the child is included as a control. The validity of this assumption is discussed in Chapter 8. She finds that private school choice has a positive, statistically significant effect on public schools productivity. For example, a public school in an area with moderately high private school choice (as opposed to moderately low private school choice) has 8th grade reading scores that are 2.7 percentile points higher. Greene and Kang (2004) use the same identification strategy for data from upstate New York school districts, but does not replicate the positive Hoxby result. It finds generally significant positive effects of private competition for some measures of school output, but little if any on measures such as the percent of students receiving high quality statewide Regents diplomas (for grades 9 to 12).

The use of the religious communities instruments means that Hoxby's positive effect is restricted to a Catholic private school effect. US studies that analyse the effect of all types of private schools have yielded mixed results. Hall and Vedder (2003) show a positive association between the proportion enrolled in private schools within Ohio school districts and public school performance. However, this positive (though small) effect is estimated in a simple cross-sectional OLS regression, with relatively few control variables. Other US studies have not successfully replicated her finding. For example, Geller et al. (2001) failed to find any effect of private school competition on public school performance in Georgia using a 10 year panel of data including 3^{rd} and 10^{th} grade scores for reading and maths. The models are estimated using lagged measures of competition in regressions and estimating first difference models.

The introduction of a nation-wide voucher scheme in Chile in 1981 produced a significant growth in private sector schooling. Hsieh and Urquiola (2003) examine the growth in competition from private schools in Chile on total pupil achievement in the

municipality. They employ a difference-in-differences approach with municipality fixed effects using data on 150 municipalities from 1982 and 1988. The outcomes of interest are changes in the municipality mean maths and language scores, average repetition rates and average years of schooling among 10-15 year olds. They are able to condition on changes in the SES of pupils in each school. Since the growth in the private sector is potentially endogenous to factors that affect changes in the outcome variables, they use IVs of urbanisation level, population size and heterogeneity in adult education in the area to predict the growth in private schooling. They find no evidence that competition from private schools improved average educational outcomes in the municipality. McEwan and Carnoy (2000) draw slightly different conclusions for the later period of 1988 to 1996 using a repeated 8-year cross-section and using a difference-in-differences to control for school unobservables. Their best estimates suggest that 15 years of competition in Chile led to modest gains of around 0.16-0.2 s.d. in test results among public schools experiencing competition from the private sector in Santiago. However, in other regions (where 75% of the population live) they find that this private school competition had slightly negative effects.

The effect of the growth in competition from private schools on lower-secondary state schools following the introduction of vouchers in Sweden is also contested. A study using pupil-level cross-sectional data found (mostly) positive relationships between the proportion of pupils in private schools in the municipality and ninth grade credit scores in maths, English and Swedish (Sandstrom and Bergstrom, 2005). Two political control IVs were used to predict the proportion in private schools in an attempt to deal with the endogeneity of the location of private schools. However, the administrative dataset used by them had poor pupil controls and no prior achievement variable. Furthermore, the independence of political control as an instrument is unclear, particularly given the very poor control variables, since it affects educational outcomes by acting a proxy for other cultural and social pathways beyond the incidence of private schooling. Using the same dataset and a slightly different IV approach, but for a sample for which a prior attainment measure at grade 6 exists, Ahlin (2003) finds no statistically significant effects of private school competition on municipality achievement. Bjorklund et al. (2004) broadly agree with Ahlin that the positive association found between independent school share and pupil outcomes is not robust to their using a repeated cross-section with municipality fixed effects. They do find a positive effect on the maths score in the population, rather than the sample, data, but it is small. Because it is in the administrative population data it is possible that the poor pupil-level controls are responsible for this positive association. A final study, using very similar data (from 1988 to 2003) to Bjorklund et al. (2004) found modest, but statistically significant effects of private school share on average municipality GPA test scores (Böhlmark and Lindahl, 2007). The study employed a difference-in-differences

estimation strategy, looking across municipalities in Sweden and comparing changes in pupil outcomes with changes in private school enrollment. This is valid provided the factors that explain supply of private schools are fixed over time within municipalities. It is difficult to understand why these final two studies disagree in terms of the effect of private schools. However, Böhlmark and Lindahl (2007) report fewer robustness checks than Bjorklund et al. (2004) so it is possible that their findings are sensitive to exact specification. It is worth making a comment on the outcome variable used by all Swedish studies of school performance. The 9^{th} grade achievement is assessed by teachers, so there is a risk that 'measurement error' on this test (i.e. differences between teacher reported versus actual competency) is correlated with the pressure the school perceives themselves to be under from the education market.

Bradley et al. (2004) provide a plausible explanation for why so few international studies find a strong positive relationship between competition and school performance. They calculate the relative efficiency of primary schools in Queensland, Australia using DEA and show that efficiency of schools shows limited correlation with league tables that do not adjust for socio-economic background of pupils and other school inputs. Thus, the incentive for schools to increase their efficiency is quite limited if league tables or raw exam results are the primary means by which parents judge schools. However, it is notable that in the UK value-added data is published, yet parents do not appear to utilise it (as discussed in Chapter 3). In Australia they do show that school efficiency is positively affected by between-government school competition, albeit weakly. However, there is weak evidence of net negative effects on efficiency of cream-skimming by non-government (voucher) schools.

6.3.6 International 'choice' school competition evidence

Tel Aviv in Israel used to operate a system of bussing to integrate pupils from rich and poor neighbourhoods. Victor Lavy (2006) examines the effect that abandoning a school bussing program and introducing choice had on the pupils of one school district in Tel Aviv. He employs a regression discontinuity design with comparison groups from untreated tangent neighbourhoods to examine the overall effect of the choice program on pupil outcomes in lower secondary schools in terms of drop-out rates, matriculation rates and average scores in matriculation exams. He finds the choice program had significant general effects on all these pupil outcomes, particularly among disadvantaged children, and attributes this to the threat of competition (which was very real because schools actually closed) and on better matching of pupils to schools.

6.4 Discussion

Overall, the international evidence on the effects of competition on pupil performance is mixed. In the UK, early studies yielded some positive associations, but the measures of competition used were highly correlated with population density and none had access to pupil-level data. However, the recent null results of Gibbons et al. (2008a) and Burgess and Slater (2006) may result from methodological considerations, particularly regarding their measurement of competition, and so should not necessarily be taken as decisive evidence of no relationship between competition and school performance. The following two chapters increase the evidence on competition between schools in England by exploiting geographical variation in the supply of autonomous – foundation and religious – state schools to estimate the effects of competition on pupil exam performance. The analysis benefits from the availability of pupil-level data and draws on identification methods used by US economists working in the same field.

The US non-experimental evidence is similarly mixed, and where consistent positive effects are estimated they are usually quite small. Certainly, the growing body of evidence from the Charter school programme could not yet be described as consistent. In any case, the majority of these papers estimate short-run effects of policy changes, and longer-run effects may turn out to be very different. Furthermore, the extent to which UK policy-makers are able to draw useful conclusions from US literature, given very significant institutional differences in schooling structures, must be limited.

The effect of private schools on public school systems across the world is once again not consistent, even where researchers are analysing exactly the same education system. It is perhaps not surprising that the effect of competition from the private sector is somewhat muted. Parental choice of private school may have little to do with the performance of public schools, particularly in countries where private schooling is predominantly religious or the preserve of a wealthy élite. Alternatively, public school competition in these areas may already be sufficiently high such that private schools provide no additional effective competition. Equally, private schools may not perform any better than public schools and thus not present a significant threat.

The large number of non-experimental studies that find no effect of competition on pupil performance raises the question of whether this is because competition is complex to measure, or competition is complex to implement, or both. In terms of implementation of policies to encourage competition, it is possible that headteachers are aware of competitive threats, but are not able to translate their increased effort into greater effort on the part of teachers. This might be the case if teachers were displaying knavish behaviour, particularly given the negligible relationship between teacher pay and performance in most countries discussed in this chapter. It would also be true of

teachers who had knight-like attitudes towards their work, but who cared about pupil outcomes beyond test scores. Alternatively, headteachers and school administrators may perceive themselves to be operating in a competitive environment, but they do not care about the small increases or falls in pupil numbers that are possible given limited spare capacity, and so a low-effort, non-competitive equilibrium persists.

Alternatively, perhaps levels of competition do vary, but they never vary enough in practice for researchers to be able to identify this variation effectively in non-experimental studies. Indeed, the only study to show that choice policies really can have a substantial effect on pupil test scores comes from the best designed intervention – the choice pilot program in Tel Aviv; it is possible that similarly designed pilots in other countries would also find a substantial effect. The important aspect of this intervention, compared to most other competition scenarios, was that the competitive threat was undoubtedly large because schools really were threatened with closure (and did close). Given that it is not clear that schools aim to grow in a competitive environment, it is possible that this threat of school closure is essential in a competitive environment to produce increased effort and focus on pupil performance.

Chapter 7

Spillover effects of the Grant-Maintained schools policy

The Grant-Maintained (GM) schools policy was a relatively radical experiment in school autonomy, with one-in-six secondary schools achieving independence from Local Education Authority (LEA) control between 1989 and 1997. So it provides one of the best opportunities to assess whether market-like structures to introduce choice and competition in the English schooling system are capable of improving educational outcomes for pupils at age 16. Evaluating the gain from school autonomy and its impact on neighbouring schools is particularly important given similarities between the GM schools policy, the ongoing academies programme and new legislation to encourage Community schools to acquire Foundation Trust status (and thus operate independently from Local Authority control).

Although the policy was suspended shortly after the Labour Party came to power in 1997, there is reason to believe that it has had a lasting impact on schools. This chapter therefore seeks to study this *long-term* impact of the GM policy using recent pupil-level data. Legislation in 1998 gave GM schools a new Foundation status (or they reverted to Voluntary-Aided status). Under their new status, they retained the same control over operations and admissions as before, gradually lost their preferential financial treatment and were required to have some LEA-appointed governors (Anderson, 2000).[1] Nevertheless, this continued relative autonomy may bring benefits to the former GM schools, and impact positively or negatively on their neighbours, even today. It is true that Foundation schools appear to out-perform Community schools in both Contextual Value-Added (CVA) and standard GCSE league tables, but it is

[1] They could also obviously continue to benefit from past superior capital investment

important to understand whether this apparent superior performance should be taken as evidence that governance structure can influence school effectiveness or whether it is biased by sorting based on unmeasured pupil characteristics.[2] Indeed, Chapter 4 showed that Foundation schools are associated with smaller proportions of pupils in the LA attending their nearest school, which suggests they are enabling some degree of parental choice and sorting, thereby increasing the competitive threat experienced by neighbouring schools. Thus, the policy provides an opportunity to look for evidence supporting the thesis that competition for pupils improves school performance across an area. It is these area-wide estimates – the combination of own-school and spillover effects – that are the relevant metric by which to judge the effect of this important policy on social welfare.

Table 7.1 summarises the GM schools policy, in relation to Community and Foundation schools today. GM schools were owned and managed by their governing bodies. They received funding directly from the Department for Education which in turn recouped the cost from the former LEA's revenue support grant. They were completely independent of the former LEA and accountable directly to the Department for Education (and later the Funding Agency for Schools).[3] The GM school governing bodies dealt with all staff matters, including suspension and dismissal, though staff enjoyed the same pay and conditions as prior to the transfer of contracts.[4] Some GM schools were allowed to introduce sixth-form provision, against the wishes of the LEA. It should be noted though that, although the GM reform does appear quite radical, it was introduced against a background of increased financial autonomy for all schools under Local Management of Schools (LMS) reform. Indeed, Deem and Wilkins (1992) argue that GM and LEA-LMS schools should not be regarded as different in kind, but as institutions located at different points on the same [self-governing] continuum.

There were three distinct claims made for the GM schools legislation at different times. First, the GM policy put in place both the incentives and capabilities for substantial improvements in school efficiency at these schools. This may have been possible because schools were given control over all financial, building and staffing decisions. However, their (mostly) higher levels of per-pupil funding make it difficult to attribute any superior performance to school autonomy rather than increased resources.[5] Sec-

[2]See, for example, Department for Education and Skills (2007a).

[3]The LEA was still required to provide certain services to pupils at opted-out schools. These included education welfare, careers, home-to-school transport, and support for statemented pupils (Bush et al., 1993).

[4]One survey indicated GM schools made greater use of fixed-term and part-time contracts (Thompson, 1992).

[5]Bush et al. (1993) find that almost all schools in survey report higher spending on books and materials, furniture, maintenance and decoration, INSET, support staff. Many report higher spending on teaching staff levels, but few increased salaries as such. Technology and science were major beneficiaries of capital works. Simon and Chitty (1993) report that GM schools got on average four times as much in the way of capital grants than mainstream county schools.

Table 7.1: Differences between Grant-Maintained, Foundation and Community status

Grant-Maintained schools (1989 onwards)	Foundation schools (1998 onwards)	Community schools
Land and buildings owned by Governing Body, who make decisions about new capital spending, funded directly by central government.	Land and buildings owned by Governing Body, who make decisions about new capital spending, funded via the formula grant.	Land and buildings owned by LA, who makes decisions about new capital spending, funded via the formula grant.
Funded directly from central government, often at significantly higher levels than equivalent LA schools.	Funded via LA under same terms as other schools, although higher levels of GM funding at schools took many years to phase out.	Funded via LA.
Managed by their governing bodies, with no accountability to the LA.	Managed by their Governing Body, with minority LA representation. LA may intervene in very specific circumstances.	Managed by their Governing Body, with majority LA representation.
Governing Body runs the selection panel, appoints and employs the headteacher.	Governing Body runs the selection panel, appoints and employs the headteacher.	LA appoints and employs the headteacher, but the selection panel is usually run by the Governing Body.
Staff are employed, appointed and dismissed by the Governing Body.	Staff are employed, appointed and dismissed by the Governing Body.	Staff are employed by the LA. The Governing Body appoints staff, but the LA has advisory rights.
Governing Body has responsibility for deciding and administering admissions arrangements, with changes needing approval by the Secretary of State. Schools may apply to introduce academic selection.	Governing Body has responsibility for deciding and administering admissions arrangements, which must adhere to the Admissions Code. All pre-existing selection may remain.	LA has responsibility for deciding and administering admissions arrangements, which must adhere to the Admissions Code.

Source: Education and Inspection Bill 2006 Annex B, accessed July 2007.
http://www.publications.parliament.uk/pa/ld200506/ldbills/116/en/06116x-k.htm

ond, politicians argued it would increase parental choice through diversity of school provision; but it is widely agreed that regulations combined with restrictions imposed by the National Curriculum prevented GM schools from pursuing a distinct curriculum or identity (Bush et al., 1993; Fitz et al., 1993b; Sherratt, 1994).[6] Finally, it was suggested that the policy would be a catalyst for improvement in standards across the system as GM schools provided a competitive threat to their neighbouring schools (or alternatively induced local authority officials to work harder to ensure other schools did not want to leave their control). Clark (2007), in a paper discussed in some depth in the following section, showed there was little short-run evidence for this competition claim.

The GM schools policy was very controversial, being opposed by LEAs of all shades; some Anglican and RC churches; some Department for Education officials; large areas of the press; and teacher unionists (Sherratt, 1994). So, the government required the Governing Body of schools wishing to acquire GM status to pass a resolution proposing that an election be held; then win a majority vote of the parents of current pupils. About two-thirds of the c.850 secondary schools who took this vote gained over 50 per cent of the parental vote and thus became GM schools.

Given high levels of opposition to the policy, the decision to take the parental vote could not be taken lightly. According to Fitz et al. (1993b), headteachers were instrumental in deciding whether a vote should be held. A survey of headteachers, chair of governors, teacher union reps, parent governors and teacher governors in Bush et al. (1993) reports that the main stated motivation for gaining GM status on the part of the school was the increased autonomy and higher levels of funding that were available (see Figure 7.1). However, a notable minority did so to avoid closure or redesignation (94 of the first 439 schools to embark on the process were subject to a Section 12 or 13 closure or redesignation proposal; many others – up to half in early surveys – were not yet subject to formal plans but believed they would be in the future (Bush et al., 1993)).[7] A final group of schools had some disagreement with LEA officials (for example, over tertiary education policy or staffing decisions). This heterogeneity in motivations to seek GM status is important to quantitative analysis of the policy because it is likely to result in heterogeneity in benefits to autonomy at different schools.

The fact that a vote of parents was required to become a GM school provides an important identification strategy for evaluating the policy experiment. We are clearly concerned that those schools taking the vote are a non-random selection of all schools in England, thus any improvement in test scores might be attributable to systematic unobservable characteristics of this group of schools. However, under assumptions

[6] Fitz et al. (1993b) report only a tiny minority of parents who felt the new GM status had any bearing on their choice of school.

[7] A specific request by an LA to the Secretary of State to either close or re-designate a school.

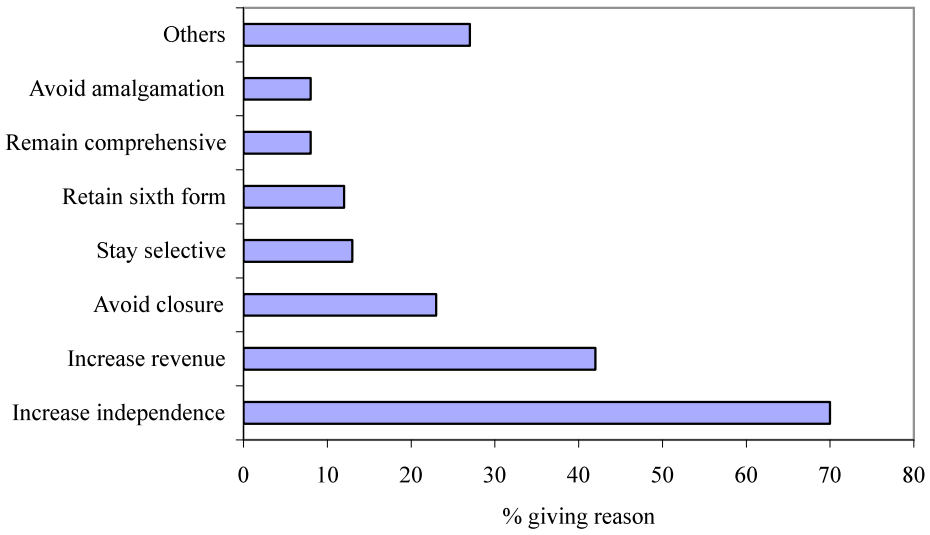

Figure 7.1: Reasons given for opting out (from Bush et al., 1993, Table 8.1)

discussed in the next section, Clark (2007) used the schools that just lost the vote to gain GM status as the policy counterfactual to those who just won the vote. He was thus able to show significant short-run gains in GCSE results for those winning the GM vote, relative to those who lost the vote.

This chapter builds directly on the work of Clark, but uses recent data to look at the long-term effects of the policy. The next section examines the validity of the regression discontinuity as a means of evaluating the impact of the GM schools policy. The validity of the empirical work in this chapter is contingent on some of the validity tests contained in Clark, and also additional new tests, so they are reviewed here. Section 3 looks for evidence that the GM policy had 'spillover' effects on neighbouring schools. It presents causal evidence on whether former GM schools may be cream-skimming more affluent or higher ability pupils. The fourth section also looks at spillover effects by measuring the average GCSE performance of pupils in areas with former GM schools to look for evidence of competition effects. Finally, the chapter aims to decompose the area effects by looking at whether (former GM) Foundation schools are more effective as a result of their greater operational independence.

7.1 The use of the GM schools vote in a regression discontinuity design

A regression discontinuity design (RDD) allows us to identify the effect of becoming a GM school using the schools that lost their parental vote as the control group. The

dichotomous treatment, d, is a deterministic function of a single, observed continuous covariate, V_j, the percentage of parents voting yes to GM status in school j. Treatment is assigned to those schools whose vote share is greater than 50 per cent ($d = 1\,[V \geq 50]$). We cannot estimate $E[y_1 - y_0]$ as $E[y_1 \,|\, d = 1] - E[y_0 \,|\, d = 0]$ because we suspect that a set of (observed or unobserved) covariates, x,[8] alter both the probability of receiving the treatment *and* the proportion of pupils gaining 5 or more A*-C at GCSE, y:

$$\begin{aligned} y_j &= \beta_d d_j + g(x_j) + \varepsilon_j \qquad d = 1\,[V \geq 50] \\ V &= f(x_j) + u_j \end{aligned}$$

The RDD used assumes subjects near the threshold of 50 per cent are likely to be similar and thus comparable. The group of schools defined as being 'near the threshold' varies through the paper from as narrow as $V = [45, 55]$ to as wide as $V = [15, 85]$, as discussed in detail later. This 'threshold' randomization identifies $E[y_1 - y_0 \,|\, V \simeq 50]$. Hahn et al. (2001) formally establish minimal continuity assumptions for identifying treatment effects in the RDD:

$$E[y_0 \,|\, V = v] \text{ and } E[y_1 \,|\, V = v] \text{ are continuous in } v \text{ at } 50 \tag{7.1}$$

The RDD estimates a weighted average treatment effect for the entire population, where the weights are the probability that the school draws a V near 50 per cent (Lee, 2005a). This means we can infer little about the potential effects of GM status for those schools who achieved very low or high V, e.g. 10% or 90%, and indeed for those who did not take the vote at all.

The GM schools' parental vote represents an unusual application of the RDD because schools have some influence over their vote – there is non-random self-selection. Specifically, the vote share V obtained by the school will be dependent on the headteacher's persuasiveness and campaigning effort (and even parents' perception of the benefit of the treatment), so that, on average, those who receive the treatment of winning the vote ($V > 50$) could be systematically more talented or ambitious than those who lost their vote ($V < 50$). Lee (2005a) explores the implications of this and shows that provided there is some random chance error component to V that has a continuous probability density function, treatment status in a neighbourhood of $V = 50$ is statistically randomised.

McCrary (2007) also argues that the continuity assumption is not sufficient where agents are able to manipulate the assignment variable, as is the case with headteachers

[8] For example, the affluence or political persuasion of parents of children at the school.

and vote share. He develops a test of manipulation related to continuity of the running variable (V) density function. The important insight is that we should not be able to see significant 'bunching' of observations that just pass the assignment threshold. Figure 7.2 shows that for the GM schools we do not see an unexpectedly large number of schools achieving vote shares between 50 and 55 per cent.

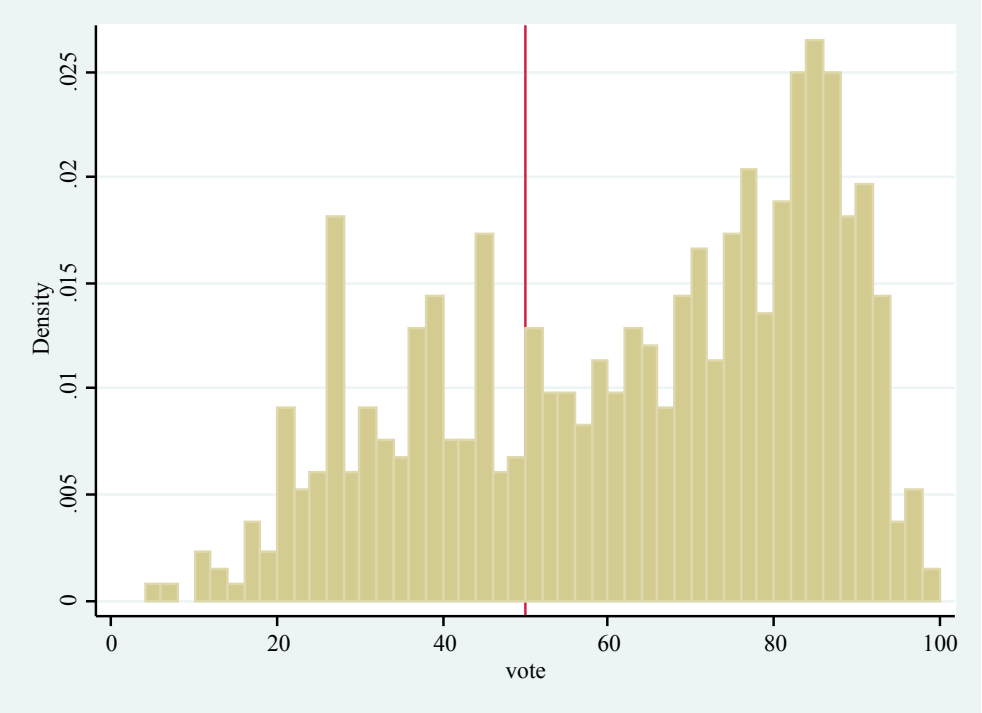

Figure 7.2: McClary's test of manipulation of assignment variable for GM voting data

7.1.1 Use of the GM vote in Clark (2007)

Clark (2007) applies the RDD to estimate the effect of winning the GM vote on school GCSE performance for between one and eight years after the change in school status. School GCSE performance is measured as the percentage of pupils gaining 5 or more GCSEs at grades A*-C at age 16 – the only available metric in datasets available for the relevant time period. Clark finds a positive effect of GM status in the order of a one quarter of a standard deviation (or 4-6 percentage points) change in school performance after two years. Figure 7.3 shows a running smoothed mean regression-adjusted pass rates from Clark (2007, Figure 3b) for 520 non-grammar schools in the wide vote share interval $V = [15, 85]$, estimated in the following form:[9]

$$y_1 = \beta_0 + \beta_1 win + \beta_2 win * vote + \beta_3 lose * vote + \beta_4 y_0 + \beta_5 controls + \varepsilon \qquad (7.2)$$

[9]The main result is estimated using 2SLS, as discussed later in this section. Controls are school type and year-term of vote.

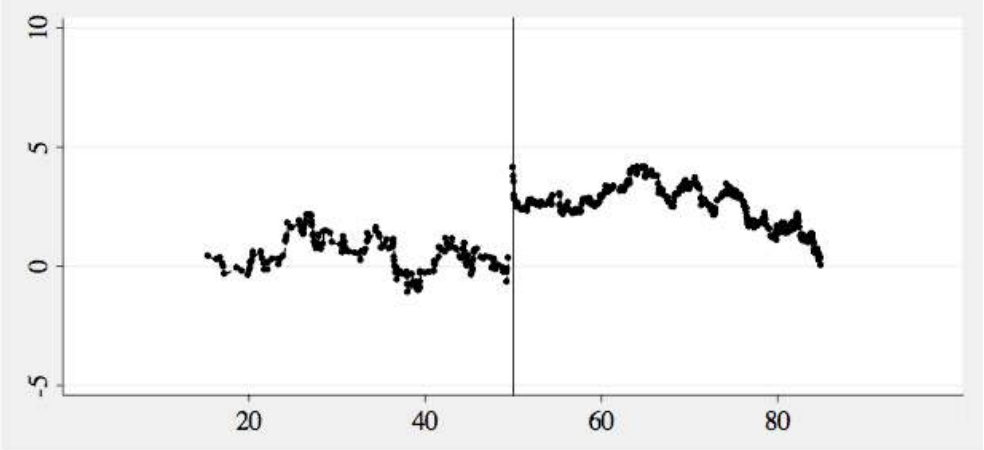

Figure 7.3: Clark's effect of a GM vote win on pass rates: 2 years after vote

Clark finds this positive effect of GM status to be persistent over the eight years following the change in school status, but is not able to control for changes in sorting (either through school exclusions[10] or more importantly through school admissions) due to a lack of pupil-level administrative data for this period.[11]

The lack of a set of school composition controls is likely to explain why Clark's findings contradict those of Levačić and Hardman (1999) who do not employ an RDD, but instead carry out a difference-in-difference analysis of the change in the performance of GM schools compared to LEA schools from 1991 to 1996. They agree with Clark that on a straightforward comparison of GCSE examination performance, the rate of improvement was higher for GM schools. However, when they added school control variables they found that this apparently superior performance could be attributed to having falling proportions of social disadvantaged students in GM schools.

The RDD shown above is one of many model specifications set out in Clark (2007) and others will be mentioned throughout this chapter. The remainder of this section discusses validity and interpretation issues that relate directly to the unusual nature of the RDD using GM schools data.

[10]There was a very sharp increase (as much as three-fold) in the number of permanent exclusions made by schools in the early 1990s, which Gillborn (1996) attributes to increased school competition and the publication of league tables.

[11]Pupil-level data for approximately one-in-ten schools where pupils took A levels between 1994 and 1997 was used in Gray et al. (2001) and is available in electronic format. This dataset reveals that GM schools with sixth forms made no improvement in A level performance over this period, controlling for GCSE attainment, gender and age of pupils. The dataset can also be used to show that this subsample of GM schools did not improve their average or total GCSE score, relative to LEA schools, for these pupils who continued to A level at the same school (and would have completed their GCSEs between 1992 and 1995).

7.1.2 RDD – efficiency versus bias trade-offs

Clark estimates the effect of winning the GM vote versus losing over a wide band of vote share values $V = [15, 85]$. This was necessary since there were too few schools close to the threshold to achieve statistical significance at the 95% level, given likely effect sizes. The further this band widens from the $V = 50$ discontinuity, the higher the likelihood that results are biased by unobservable characteristics underlying the opt-out decision. There are three quite separate sets of characteristics that are unmeasured in the education function and that V may therefore proxy. First, V reflects the effort put in by the headteacher and Governing Body to win the vote and it is possible, for example, that heads highly motivated to become GM are also highly motivated by league table position. Second, V reflects the political attitudes of the parental body, and this is correlated with socio-economic background and therefore the academic performance of their children. It could also reflect the degree of confidence that parents have in the school's headteacher more generally, or even their perceived belief in the capacity of the school to benefit from the treatment. Third, V reflects the external circumstances the school faced at the time, in particular whether or not it expected to be closed or re-organised in the near future.

Clark's approach to dealing with unobservable differences in x as we move further from the $V = 50$ discontinuity is to use linear approximations to generate simple estimates of the discontinuity gap. The dependent variable is regressed on the vote share, separately on each side of the threshold (see equation 7.2). This parametric form, which is common in the RDD literature, exploits more data than the use of a narrow band, and can therefore be more efficient. It is also possible that it generates less biased estimates of the true conditional expectation function at the $V = 50$ threshold than a simple difference in means on a narrower band, where the *true* function has a non-zero slope.

However, the critical assumption is that the parametric regression function used for extrapolation is correctly specified (Lee, 2005b). In our case, we have no *a priori* evidence that V should be a linear function of exam score growth. Indeed, this particular regression discontinuity is particularly unusual because it is not entirely clear how the assignment variable V should enter the education production function at all: we are quite vague about the x_j it proxies. The pre-test using data prior to the treatment can partially help indicate an appropriate parameterisation by showing the relationship between outcomes and vote share prior to the intervention. However it cannot reveal heterogeneities in capacity to benefit from treatment (as are implicit in Clark's specification). Figure 7.3 shows that Clark's estimated coefficients reflect an assumption of greater benefit from treatment at the discontinuity than for schools with large vote shares. This type of heterogeneity in benefits from the treatment is perfectly plausible,

but we not have a clear explanation as to why they might exist (as opposed to any other specification of heterogeneity, for example). It is important to stress that we cannot make inferences away from the 50 per cent discontinuity using an RDD, so we have no reason to believe (and indeed do not need to believe) that the estimated slopes are correct. However, by drawing on data further from the discontinuity for efficiency reasons, we must believe that the parametric form we have chosen for the data is the correct one (hence, we must believe in the plausibility of larger treatment effects for smaller V).

We should therefore be cautious in making statistical inferences from parametric regressions. On the one hand, if the polynomials (in this case first-order) are 'correct', the estimator is efficiently using data that are both close to and far from the discontinuity. On the other hand, if the true functions do not belong to the class of polynomials we select, the discontinuity will in general be biased, and may lead to erroneous inferences of statistical significance (DiNardo and Lee, 2004).

As an alternative, non-parametric procedures are also available to estimate the conditional expectation function at $V = 50$ (Lee, 2005b). These are not heavily model dependent, simply assuming:

$$y_j = \beta_d d_j + h(V_j) + \omega_{dj} \qquad d = 1\,[V \geq 50]$$

$$(1)\ \lim_{V \to 50} E[\omega_0 \,|\, V] = \lim_{V \to 50} E[\omega_1 \,|\, V]$$

$$(2)\ h(.) \text{ is an unknown function continuous at } V = 50$$

β_d is identified with the difference between right and left limits of $E[y \,|\, V]$ at $V = 50$:

$$\beta_d = \lim_{V \downarrow 50} E[y \,|\, V] - \lim_{V \uparrow 50} E[y \,|\, V]$$

Unfortunately, identifying statistically significant results using the right and left limits is contingent on having a great deal of observations very close to the discontinuity; just 60 schools had first vote outcomes between 45 and 55 per cent. This chapter will report a variety of parametric forms, as in Clark, but will also relax identification assumptions needed by estimating differences in mean outcomes for schools in a tight band around the discontinuity. It is able to do this, where Clark could not, because the availability of pupil-level data removes most of the variance in GCSE scores, thus reducing standard errors on estimates.

7.1.3 RDD – role of the pre-test

Data collected prior to the intervention allows us to partially test the internal validity of the RDD. First, we can partially test the Hahn et al. (2001) minimal continuity assumption by checking $E[y_0 \mid V = v]$ is continuous in v at 50 (see equation 7.1). This is passed in Clark's data.

Lee (2005a), in his paper on voting data and RDD, also sets out the local independence result, which implies that baseline characteristics in treatment and control must be identical for the RDD to be valid. Table 7.2 shows mean GCSE scores prior to treatment for vote winners versus losers in Clark's dataset (obtained directly from Clark and re-analysed for this study). The data does pass this pre-test at the 5% level on the $V = [15, 85]$ interval chosen by Clark for his main results. However, the difference between means is consistently negative across all possible chosen groups, i.e. vote winners had *lower* GCSE performance prior to treatment than vote losers. Interestingly, the difference is very large – over 3 percentage points – for schools very close to the discontinuity and this might explain the large growth in GCSE scores post-treatment for schools just winning the vote.

Table 7.2: T-test of difference between means on pre-treatment GCSE scores

	Vote share interval				
	[0,100]	[15,85]	[25,75]	[35,65]	[45,55]
Number of schools	662	524	357	207	62
Win	38.05	39.36	40.72	41.04	42.35
Lose	40.96	41.48	41.57	42.78	45.50
Diff (GCSE percentage points)	−2.91	−2.12	−0.85	−1.75	−3.15
T-value on Diff	2.21	1.55	0.53	0.81	0.84
Significant difference at 5%?	Yes	No	No	No	No

Note: grammar schools are excluded

This is somewhat surprising since intuitively we assume that vote winners were more likely to be in affluent areas without political opposition to GM status. However, as discussed earlier, a significant proportion of vote winners were blighted by closure or re-organisation threat. Fitz et al. (1993a) show that prior to 1991, 80% of schools gaining GM status were motivated by avoiding reorganisation; this figure fell for later votes but they estimate it to be around 50% for their sample as a whole. It is likely that this subset of schools has no/few counterfactuals in the control group. This quote from a pupil might explain not only why these schools had relatively low GCSE performance prior to treatment, but also why they benefited so much from GM status in terms of general morale, and recruitment of both pupils and teachers to the school:

> Before, less people were coming here in case it closed. Now they know it isn't ... People aren't looking at us ... [and thinking] we're going to shut

down.

(Fitz et al., 1993a, page 78)

All this means that there is an outstanding question about how the pre-treatment outcomes data should be used. We can estimate the RDD in two ways – first without using GCSE performance before treatment (y_0) as a control and second by adding y_0 as a control. If the pre-test proposed by Lee (2005a) is passed the point estimate on the effect of winning should not really change: all the y_0 control does is explain some of the variance in y_1 and thus reduces standard errors.

Table 7.3 presents some regression results from Clark's data (after two years). We substantially change our inferences by including y_0 as a control: the raw differences in outcomes two years after treatment are just 0.2 percentage points, yet our estimated effect of GM school status rises to 2 percentage points by adding the y_0 control. In other words, the borderline pass-fail of the pre-test is potentially inflating win-lose differences in $y_1 - y_0$.

Table 7.3: Sensitivity of Clark's main (ITT) estimates to changes in control variables

	No controls b (s.e.)	y_0 included as control b (s.e.)
Win	0.202 (1.394)	1.963 (0.681)
N(schools)	520	520

Note: data taken from Table 1a in Clark, 2007

7.1.4 RDD – independent causal effect of V on y

In identifying the causal effect of GM status on school GCSE performance, we must assume that the random draw of V does not itself have an impact on the outcome, except through its impact on treatment status (Lee, 2005a). That is, while V is allowed to be correlated with y_1 or y_0 in the population, V is not permitted to have an independent causal impact on y for a given school.

There is a plausible argument that in our case V does have an independent causal impact on y. A school that wins its GM vote may experience a 'euphoria effect' that temporarily increases staff motivation, resulting in effort directed at improving test scores. Alternatively, a headteacher who wins a controversial vote might experience an increase in respect from staff, which allows him/her to unite teachers in pursuing exam-orientated goals. Similarly, the school management who loses their GM vote may well perceive the lack of support for their proposal as a vote of confidence in the school more widely. This would be de-motivating, and may even cause some vote-losing

headteachers to leave their job.[12]

Figure 7.4 shows that ideally we wish to identify the effect of choosing to take the GM vote and winning it ($Y_w(V)$), versus those schools not taking the vote at all ($Y_n(V)$). However, we are only able to estimate the effect of taking the vote and winning versus taking the vote and losing ($Y_w(V)$-$Y_l(V)$). This might be an overestimate of $Y_w(V)$-$Y_n(V)$ if it includes the de-motivational effects of losing a parental vote.

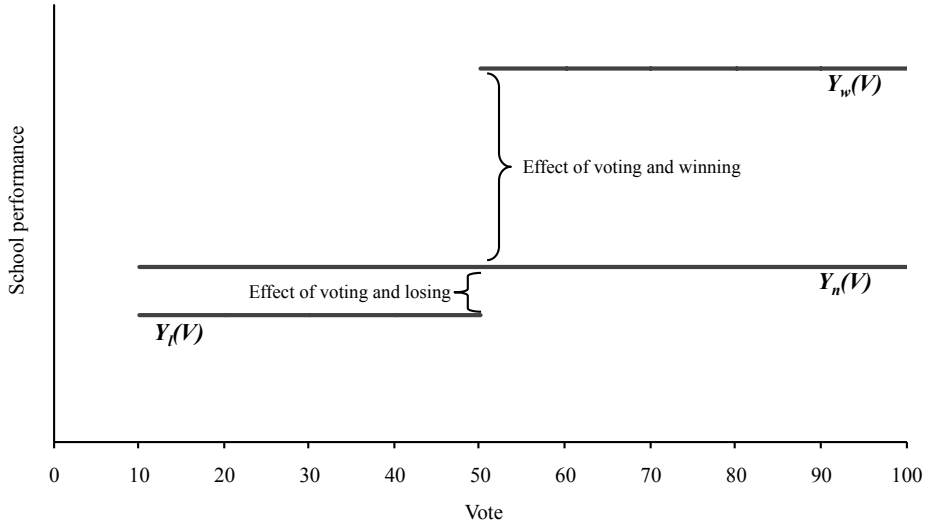

Figure 7.4: Illustration of the separate effects of voting and the voting outcome

There is some evidence for this latter phenomenon in the aggregated GCSE performance data for the period 1992-1995 (1992/3 is the time period when the majority of schools gained their GM status). The 1992-5 increase in the proportion of pupils getting 5 A*-C at GCSE is 5.20 percentage points across all schools (including grammars). For vote winners (including grammars), this figure is 6.56 percentage points; so schools gaining GM status outperformed the average. For vote losers, this figure is 4.87 percentage points; so vote losers did worse than the average school over this time period.[13]

[12]No analysis of the headteacher turnover data has been carried out for this period, but conversations with government officials and academics associated with early research into GM votes suggests this happened on more than one occasion.

[13]However, data presented by Clark in web appendix Table 3 does not identify these differences between non-voters and vote losers for the longer time period of 1992 to 2001. It is not clear why our observations disagree.

7.1.5 RDD – fuzzy versus sharp discontinuity

There are further issues in the GM votes data concerning schools that lost their first vote to become GM, but that went on to hold subsequent votes, which they won. We use the first ballot win as our assignment to treatment variable since we think it is this V that reveals the underlying characteristics of the school (x_j). However, it is the final ballot that determines the long-run treatment status of the school. A small number of schools winning the vote were rejected for GM status by the Secretary of State because they were subject to reorganisation or closure (Fitz et al., 1993a); a number of schools losing the first vote went on to hold further parental votes that they then won. Figure 7.5 shows that 25 of the 233 schools losing their GM vote went on to hold second or third votes of parents, which they eventually won, thus becoming GM schools.

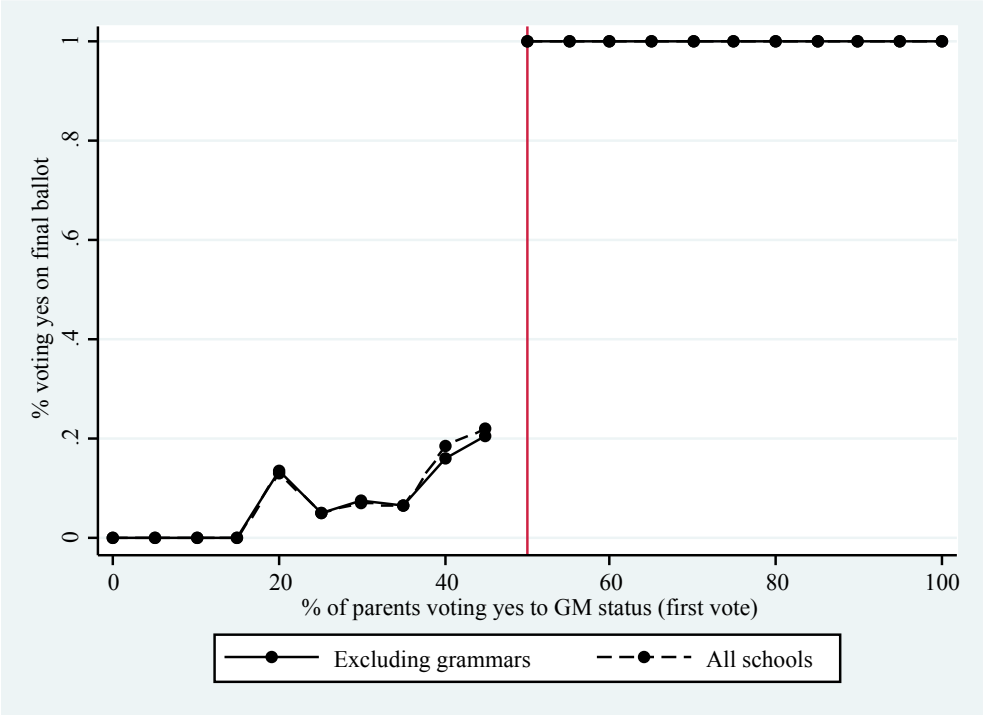

Figure 7.5: Percentage of first ballot vote losers achieving subsequent yes votes

There are several possibilities for dealing with this problem in the data. We can use the first vote data in a 'sharp' RDD, as described earlier, and interpret results as 'Intention to Treat' (ITT) estimates (Angrist et al., 1996). However, in this case it is not clear the ITT estimates are the ones we want since the GM policy is no longer in place: we are interested in the effects of school autonomy and increased competition and not the effects of taking the GM vote. Alternatively we can use the outcome of a first vote (WIN) as an instrument for the outcome of the final vote (and thus the

treatment – GM), known as a 'fuzzy' RDD in the literature:

$$y_j = \beta_0 + \beta_1 GM_j + f(V_j) + \varepsilon_j$$
$$GM_j = \alpha_0 + \alpha_1 WIN_j + f(V_j) + \upsilon_j$$

This identifies the Local Average Treatment Effect (LATE), which in this case would be the effect of receiving the treatment for schools who won the first vote (*and* had an expected vote share close to the discontinuity). As in Clark, both the ITT and IV estimates are reported; given so few schools held a second vote, the coefficients are generally quite similar. Of course, if treatment effects are heterogeneous, which they are likely to be since schools' motivations for taking the vote were very different, average effects across the whole group of GM schools might be very different; this study can say nothing about them.

7.1.6 Using RDD in recent pupil-level data

There are difficulties with using data over 15 years after a policy is introduced to identify long-term effects. First, sorting of pupils across schools, neighbourhoods and even cities will have occurred in this time period. Indeed, theory predicts this is a likely outcome of the intervention. This makes it more difficult to distinguish between school effectiveness and unobserved pupil characteristics for the 2005 cohort of school leavers, compared to pupils in the first five years of the GM intervention who entered school prior to the change in status. However, pupil data not available in the 1990s gives us information on where pupils in schools from 2002 onwards live, relative to where they attend schools. This allows us to qualitatively describe the nature of the sorting effects that control over admissions has on the allocation of pupils across schools.

On the other hand there are some significant advantages to using an RDD design in the recent pupil-level data, compared to the historical school performance tables used by Clark. The first advantage is that we can measure pupil achievement using the capped GCSE point score on their best eight subjects, a measure not available to researchers in the 1990s. The percentage gaining 5+A*-C at GCSE is a rather crude measure of achievement that can be manipulated by schools. This is why, after the publication of league tables using this measure in 1992, some schools were able to rapidly improve their league table position with little extra effort by focussing solely on grade C/D borderline pupils and offering 'easier to teach' subjects (Wilson, 2003). Capped GCSE point score is a better reflection of school performance because it captures the performance of all pupils across a wide variety of subject areas.

The second advantage is that we can be less concerned about the short-run negative

(and positive) effects of losing (and winning) the GM vote confounding our estimates of *treatment* effect since, for example, negative effects of losing a vote on school morale and headteacher turnover should have disappeared. Thus, any differences between vote winners and losers at the discontinuity should be attributable solely to the former GM schools' continuing control over admissions and management of resources and staffing (as Foundation or Voluntary-Aided schools).

A final advantage is that models can be estimated using pupil level data with a wide set of prior attainment and socio-demographic controls (described in the next section). This approach to estimating a RDD with controls is formalised in a non-parametric form by Frölich (2007), who shows it requires fewer assumptions than the one-dimensional nonparametric regression analysed by Hahn et al. (2001). The use of pupil-level data can introduce a large set of covariates to account for differences around the threshold, thus reducing bias on estimates of school effectiveness (and without dimensionality problems since the number of pupils is very large). In addition, accounting for covariates will also substantially reduce the variance in the outcome GCSE variable, thus increasing the precision of estimates. A consequence of this is that statistical significance can be achieved for the relatively narrow band of vote winners and losers.

7.2 Data

Data for the 573,227 school-leavers is drawn from the 2005 National Pupil Database (NPD). Table 7.4 summarises the key variables. The main outcome variable used is pupil achievement in GCSE exams at age 16. This is measured as total points achieved by pupils over their best 8 GCSE/GNVQ or equivalent examinations. The GCSE exam is graded from A* (58 points) to G (16 points), with 6 point increments between grades. The typical pupil achieves around 4 Cs and 4 Ds.

There are several pupil-level variables in the dataset that make use of the pupil postcode in NPD. First, the home-school distance is calculated for each pupil, and the median value for each school is found. Second, two continuous measures of deprivation – the Index of Deprivation 2004 and the IDACI measure – are included, as are 57 ACORN dummies of household type. These are imperfect indicators of the child's social background to the extent that they measure average social characteristics of households in the postcode or SOA.

Table 7.4: Key descriptive statistics

	Won vote first time	Lost first vote	Never took vote
Capped GCSE score	312.1 (98.7)	299.0 (101.6)	283.8 (106.7)
KS2 test score over 3 subjects	4.59 (0.70)	4.51 (0.72)	4.43 (0.75)
FSM eligibility	8.87%	11.80%	14.60%
White ethnicity	83.41%	83.64%	83.95%
SEN with statement	2.06%	2.22%	2.63%
SEN without statement	11.11%	11.78%	14.17%
English not first language	4.81%	9.68%	8.81%
Home-school distance for school's median pupil (km)	3.12 (2.55)	2.40 (1.51)	2.31 (1.70)
Number of pupils	111,840	47,541	413,846
Number of schools	618	233	2,257
Won subsequent re-ballot		10.8%	

Note: statistics include grammar schools, but these are dropped for some analysis. These descriptive statistics are for a slightly larger set of than Clark uses. Some vote-taking schools closed or have been renamed, but I include these schools where a new school opens on exactly the same site or where a school moves but retains exactly the same name; I have also included schools achieving GM status before 1992.

7.2.1 School variables

The voting data for these schools is taken from the Grant Maintained Schools Database, which provides details of all GM status ballots taken by schools, and the outcomes of these ballots. Further details can be found in Chapter 2.

The 3,108 schools in the dataset are identified by whether they are a grammar, Foundation (non-grammar), Voluntary-Aided (VA) (non-grammar), Voluntary- Controlled (VC) (non-grammar), academy or City Technology College (CTC). The default school is a Community (LEA controlled) comprehensive school. Foundation schools are overwhelmingly former GM schools, as shown in Table 7.5.

Table 7.5: Governance of GM vote winners and losers

	Final GM ballot won	Final GM ballot lost	GM ballot never taken
Grammar	94	8	70
CTC/Academy	0	0	14
Community (non-grammar)	12	162	1750
Foundation (non-grammar)	404	1	22
Voluntary-aided (non-grammar)	133	26	331
Voluntary-controlled (non-grammar)	0	11	69
Total	643	208	2256

For most of the regressions, grammar schools are excluded. This is because a very large majority of them won their vote, so there is not a large enough counterfactual to identify the effects of becoming a GM grammar school.

School peer group control variables used in regressions are the percentage of girls in the school; the mean KS2 prior attainment of intake; the percentage of pupils who are FSM eligible; the mean deprivation value for pupils.

7.3 GM schools and cream-skimming

Chapter 5 showed that own admissions schools – Foundation and Voluntary-Aided schools – showed signs of cream-skimming in the sense that they had intakes of higher ability and lower FSM than if they educated the pupils who lived closest to the school. This evidence was far weaker for Foundation schools than it was for VA schools. Nevertheless, it is consistent with West et al. (2004) whose survey data of 2001 showed both Foundation and VA schools were more likely to have admissions policies designed to favour easier to teach pupils and exclude others.

This section asks whether the current cream-skimming patterns of Foundation schools can be directly attributed to their having been given GM status in the 1990s. An alternative hypothesis, for example, would be that schools who were motivated to maximise their league table position (and thus cream-skim) were more likely to apply to become GM schools. Two different types of analysis are shown here. First, the pupil intake characteristics of vote winners and losers are directly compared. If they differ at the discontinuity it could be for one of two reasons: either more affluent families have been attracted to live close to vote winning schools; or these schools have used their control over admissions to systematically favour certain pupils. To distinguish between these two explanations, the second part of this section compares the intakes of vote winners and losers to the characteristics of families located close to these schools.

Figure 7.6 shows the current average KS2 point score of pupils at schools who took a GM vote in the 1990s, plotted against their vote share in the first vote of parents (rounded down to nearest 5%). If we exclude grammar schools, the evidence for GM vote winners having a superior intake is present, though weak.

Table 7.6 shows output from regressions of the same data. There are six different specifications shown here. The first four columns of estimates compare mean KS2 prior attainment of intake for non-grammar schools, varying the subset of schools on which the regressions are estimated each time. On each occasion, the parametric assumption is that KS2 intake should not vary by vote share on each side of the discontinuity for the subset of schools on which the model is estimated. This is clearly a less onerous

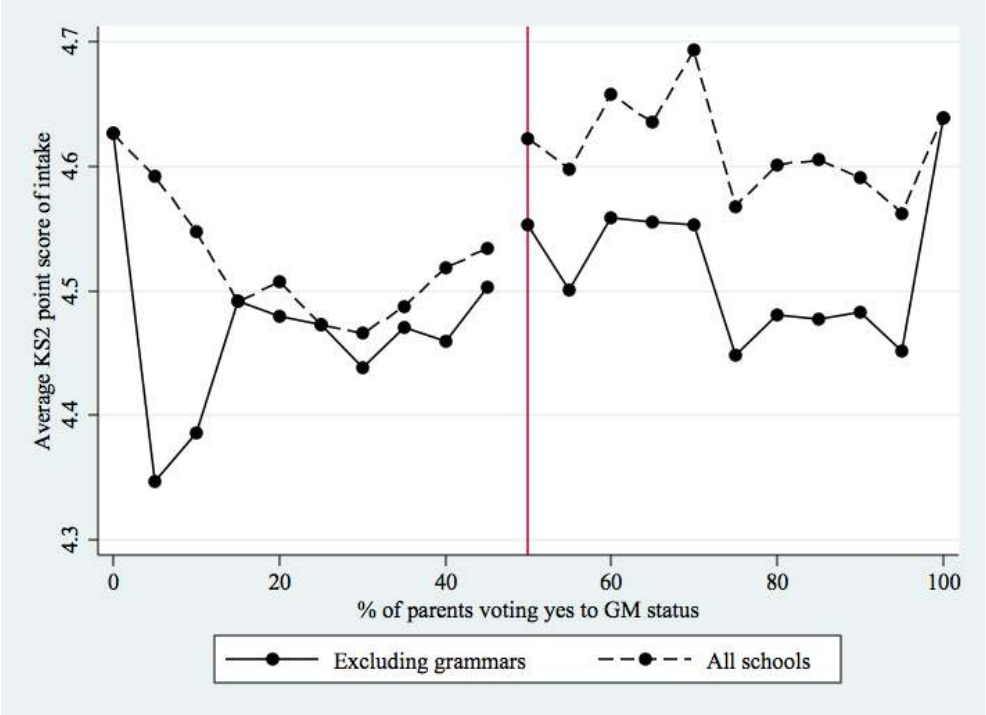

Figure 7.6: Current KS2 prior attainment of intakes for GM vote takers

assumption for $V = [45, 55]$. The tighter bands around the discontinuity reduce the chance of mis-specification biasing results, but increase inefficiency (and therefore type II errors) since sample size is small. The final two columns allow the effect of the vote on KS2 intake to vary, as Clark did in his preferred specifications. However, they rely on a very wide bandwidth. Both the ITT coefficient (whether the school won the first vote) and the two-stage IV coefficient (instrumenting GM school status using whether the first vote was won) are reported for each specification.

Overall, the evidence for cream-skimming here is present but quite weak: coefficients are always positively signed, but not always significant at the 5% level so we cannot always reject the null of no ability cream-skimming taking place. If there is a cream-skimming effect, this data suggests it is quite small: a coefficient of 0.05 corresponds to 5% of a Key Stage improvement (or 7% of a S.D.) in mean KS2 test score of entrants to the school.

Turning to evidence of cream-skimming based on social status, as measured by FSM eligibility, the evidence here exists but is also quite weak. Figure 7.7 shows that vote winners do have slightly fewer FSM eligible pupils at their schools. There is a 2.4 percentage point difference in the FSM proportions for non-grammar vote winners and losers (9.7% versus 12.1%). This magnitude of difference is consistent with Clark's data from a slightly earlier period.

Table 7.6: Differences in intakes by KS2 for non-grammar schools taking GM vote

	45%–55%	35%–65%	25%–75%	15%–85%	15%–85% with linear term	15%–85% with split linear term
Win (IV)	0.048	0.043	0.060**	0.030	0.094**	0.077
	(0.082)	(0.037)	(0.025)	(0.022)	(0.047)	(0.051)
Vote					0.156	
					(0.099)	
Vote*Win						0.238**
						(0.108)
Vote*Lose						−0.040
						(0.194)
Cons	4.515***	4.498***	4.485***	4.488***	4.455***	4.489***
	(0.070)	(0.029)	(0.020)	(0.018)	(0.028)	(0.042)
Win (ITT)	0.034	0.036	0.053**	0.027	0.078**	0.062
	(0.058)	(0.031)	(0.022)	(0.019)	(0.038)	(0.040)
N(pupils)	11,263	43,574	75,906	115,107	115,107	115,107
N(schools)	60	221	384	593	593	593

Note: *** = sig. at 1%; ** = sig. at 5%.

Dependent variable is mean KS2 score. Robust S.E.s, clustered for school

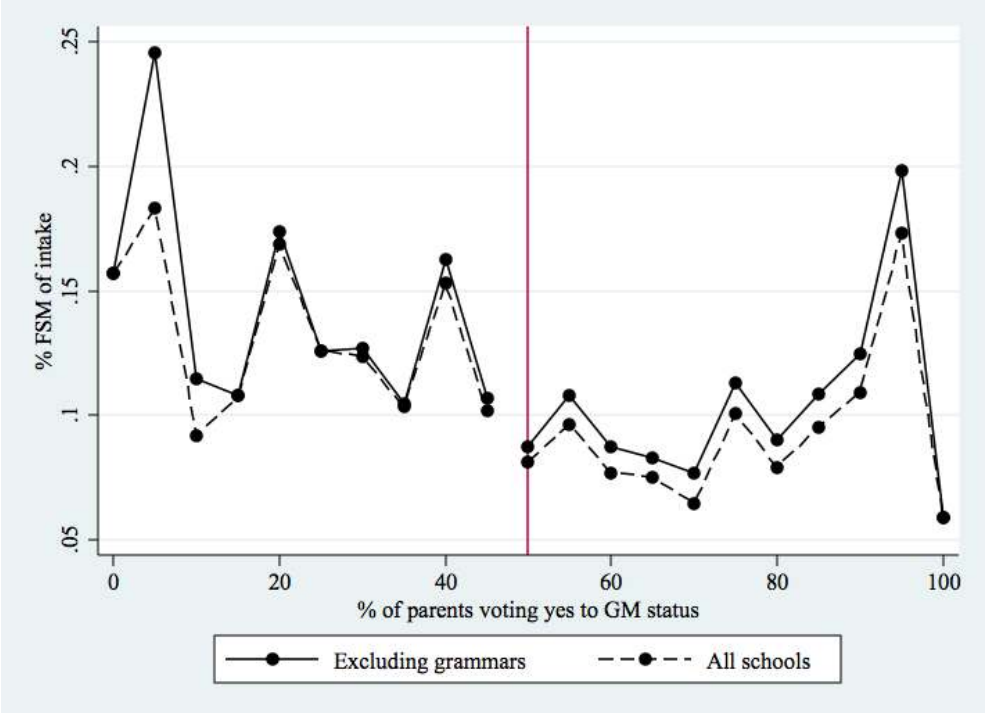

Figure 7.7: Free school meals status of intake for GM vote takers

Table 7.7 shows the estimates from a probit regression of the probability of a FSM pupil attending a vote winning versus a vote losing schools. The Win coefficient is always correctly signed – GM vote winning schools do have fewer FSM pupils – but it is not always statistically significant. A coefficient of -0.15 would indicate that a child selected at random from a former GM school has 5 percentage points lower chance that s/he is eligible for FSM, compared to a child from a vote losing school.

Table 7.7: Differences in intakes by FSM for non-grammar schools taking GM vote

	45%–55%	35%–65%	25%–75%	15%–85%	15%–85% with linear term	15%–85% with split linear term
Win (IV)	−0.026	−0.103	−0.164***	−0.172***	−0.118	−0.090
	(0.214)	(0.086)	(0.060)	(0.052)	(0.113)	(0.128)
Vote					0.131	
					(0.245)	
Vote*Win						−0.062
						(0.257)
Vote*Lose						0.518
						(0.508)
Cons	−1.336***	−1.230***	−1.204***	−1.167***	−1.194***	−1.261***
	(0.200)	(0.067)	(0.050)	(0.044)	(0.068)	(0.107)
Win (ITT)	−0.025	−0.087	−0.146***	−0.153***	−0.098	−0.073
	(0.157)	(0.072)	(0.053)	(0.046)	(0.094)	(0.102)
N(pupils)	11,506	44,755	78,014	118,529	118,529	118,529
N(schools)	60	221	384	593	593	593

Note: *** = sig. at 1%; ** = sig. at 5%.

Dependent variable is mean KS2 score. Robust S.E.s, clustered for school

If former GM schools do have slightly superior intakes today, we want to be able to distinguish whether this results from more affluent families moving closer to these former GM schools or whether the schools appear to have admissions policies or procedures that are different to those of the vote-losing schools.

We can calculate the social and ability composition of a school compared to its neighbourhood using a proximity ratio of the current composition of the school over the characteristics of the same number of pupils who live closest to the school, the construction of which is described in Chapter 5. These ratios are created for the proportion of FSM pupils in the school, the proportion of pupils who scored in the top quintile on KS2 tests and the proportion of pupils who scored in the lowest quintile on KS2 tests. Regressions are run with these ratios as the dependent variable for GM vote-taking schools. Table 7.8 shows the output for these regressions where vote winners are compared to vote losers (and the win coefficient is instrumented using the first vote outcome). All the regressions display coefficients on the win variable that are

signed as expected – GM vote winners have intakes that are more superior to their local neighbourhood (compared to GM vote losers). However, they are not always significant at the 5% level.[14]

Table 7.8: Ratio of current-to-proximity intake characteristics

		45%–55% b (s.e.)	35%–65% b (s.e.)	25%–75% b (s.e.)	15%–85% b (s.e.)
	Win (IV)	−0.129	−0.235*	−0.241***	−0.203***
		(0.186)	(0.122)	(0.077)	(0.069)
FSM	Cons	1.039***	1.234***	1.190***	1.201***
		(0.148)	(0.090)	(0.058)	(0.056)
	Adj R-Sq	0.000	0.019	0.022	0.014
	Win (IV)	0.344**	0.173*	0.060	0.048
Top 20%		(0.172)	(0.091)	(0.111)	(0.081)
KS2	Cons	0.870***	0.995***	1.090***	1.066***
score		(0.136)	(0.067)	(0.083)	(0.065)
	Adj R-Sq	0.035	0.025	0.000	0.000
	Win (IV)	−0.042	−0.073	−0.120**	−0.103**
Low 20%		(0.122)	(0.075)	(0.050)	(0.041)
KS2	Cons	0.976***	1.060***	1.078***	1.075***
score		(0.097)	(0.055)	(0.038)	(0.033)
	Adj R-Sq	0.000	0.006	0.014	0.012
	N(schools)	58	215	377	581

Note: *** = sig. at 1%; ** = sig. at 5%.

GM schools impact on their neighbours, albeit weakly, by cream-skimming more affluent pupils. Furthermore, Table 7.9 shows that GM schools are likely to have presented some threat to their neighbouring schools, because the policy appears to have allowed them to grow, relative to vote losers. The effect size here is not consistent across specifications, but the magnitude of their growth in size appears to be in the order of 10% (consistent with Clark).

7.4 GM schools and competition

This section looks at the effect of GM schools on the GCSE outcomes of pupils attending schools in the wider area. The previous section gives evidence that former GM are responsible for some sorting of more affluent pupils into their schools. Chapter 4 shows that Foundation schools are associated with a smaller proportion of pupils attending their nearest schools. All this suggests that former GM schools must be actively competing with their neighbours for pupils, and so this competition *may* increase area-wide

[14]Regressions of proximity intakes of vote winners and losers show there are no differences, suggesting families have not moved to live closer to GM schools.

Table 7.9: Differences in school size between vote winners and losers

	45%–55%	35%–65%	25%–75%	15%–85%	15%–85% with linear term	15%–85% with split linear term
Win (IV)	148	79	82 **	−11	199 ***	211 ***
	(114)	(53)	(41)	(34)	(72)	(77)
Vote					515 ***	
					(154)	
Vote*Win						457 ***
						(174)
Vote*Lose						660 **
						(293)
Cons	1011 ***	1107 ***	1111 ***	1148 ***	1044 ***	1020 ***
	(91)	(39)	(31)	(28)	(42)	(60)
Win (ITT)	110	68	73 **	−10	168 ***	174 ***
	(85)	(46)	(36)	(31)	(61)	(63)
Adj R-sq	0.04	0.02	0.02	0.00	0.02	0.02
N(schools)	60	221	384	593	593	593

Note: *** = sig. at 1%; ** = sig. at 5%.

pupil achievement.

Clark argues that there are three possible routes by which schools that neighbour GM schools might improve their GCSE results. First, there might be a direct competition effect, whereby increased competition for (certain) pupils induces neighbouring schools to increase effort directed at exams (thereby improving league table position and making the school more attractive to parents). Second, the presence of successful GM schools might have induced neighbouring schools to also pursue autonomy from the local authority (although he finds no evidence of this 'copycat' behaviour). Third, the presence of GM schools might encourage LA officials to work harder in the interest of remaining schools to discourage them from pursuing autonomy (this would seem to be more relevant to estimating short-run effects). Clark argues that GM schools could give other local non-GM schools '*an important bargaining chip: either local authorities accede to their demands (for example, to fire certain teachers or suspend certain students) or those schools would themselves seek GM status*' (Clark, 2007, page 5). However, it is of course entirely possible that this particular threat to LAs produced a negative impact on pupil test scores since LA officials might be more reluctant to intervene in schools where there were problems.

This encouragement of direct competition between schools was the reason why the GM policy was opposed by many headteachers, who feared the effects it would have on the dynamics of schooling in their area. Brown and Baker's survey (1992) revealed that

most of the heads who had decided not to opt out thought that GM status runs against the 'spirit of cooperation', and that their LEAs were doing the best they can given the financial limitations and other constraints within which they have to operate.

As described in Chapter 6, the possible effects of the GM policy on GCSE results of pupils living in the area as a whole are quite complex. The known cream-skimming will redistribute peer effects across the area, affecting the composition of classrooms in many schools. Table 7.9 also showed the policy caused changes in sizes of schools in the area, and if these schools are differentially effective, this will impact on average GCSE scores. Finally, there is the possibility of effort on the part of schools induced by the increased competition for pupils. What is measured in this section is the total effect of all these phenomena on average GCSE outcomes in the local authority, estimated in a pupil-level regression with capped GCSE scores as the outcome variable. Theory cannot help us predict the possible direction of the effect of GM schools because we cannot not take any a priori view as to whether any potential effects of competition are outweighed by anti-collaboration and cream-skimming effects.

In this section two LA-level variables are used to measure the extent to which the area is likely to be affected by the GM schools policy:

Area 0-100%: the percentage of pupils in the LEA currently in schools that
(a) took the GM vote;
(b) won the GM vote.

Area 15-85%: the percentage of pupils in the LEA currently in schools that
(a) took the GM vote for the $V = [15, 85]$ sub-set of schools;
(b) won the GM vote for the $V = [15, 85]$ sub-set of schools.

The *GM taken* coefficient tells us the effect of having schools that took the GM vote in the LA. The *GM won* coefficient tells us the incremental effect of having schools in the LA who both took and won the GM vote, and therefore is intended to measure the direct effects of the policy. Both coefficients are ITT so tell us the effect of winning the first vote (versus losing it).

Table 7.10 shows estimates from regressions of these LA 'competition' variables on capped GCSE point scores.[15] The five columns of estimates differ only in terms of the number of control variables that are included, except for the final column which includes selective LAs where over 10% of pupils are in grammar schools.

The fact that the coefficient on *GM won* falls from a large positive value in the first column to below zero in the third column suggests that there must be some sorting of more affluent pupils into LAs with GM schools. The preferred specification in

[15]The regression is estimated at pupil-level to control for confounding pupil characteristics (with clustering for schools).

Table 7.10 is the third column. This is because we do not want to attribute sorting of higher ability pupils into the area to the success of the policy (i.e. the first and second columns do not identify efficiency improvements). Also, we do not want to control for school peer groups (as the fourth column does) since changes in these is part of the effect we want to measure. Finally, it would seem to be less interesting to include areas with selective schooling systems since so many of these schools acquired GM status, but they will not tell us about potential competition effects in non-selective areas.

These results show that although being located in an LA with a higher proportion of former GM schools is associated with higher raw GCSE scores, there is no statistically significant effect once pupil characteristics of the LA are taken into account. This finding is consistent with Clark who finds no spillover effects of the GM policy on GCSE results in the area.

Table 7.10: Effect of former GM schools on local authority pupil performance

		GCSE point score (best 8 grades)				
		b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)
Area	GM won	15.617	10.424	−1.045	−2.379	7.944
0-100%:		(12.120)	(7.369)	(5.783)	(5.563)	(4.898)
	GM taken	4.206	3.687	1.454	2.846	−2.666
		(10.766)	(6.635)	(5.133)	(4.910)	(4.551)
	Adj R-sq	0.001	0.459	0.542	0.545	0.553
Area	GM won	24.224*	18.849**	3.508	3.238	13.825***
15-85%:		(13.305)	(8.027)	(6.425)	(6.176)	(5.347)
	GM taken	2.398	0.115	−1.313	−0.076	−5.963
		(10.990)	(6.762)	(5.338)	(5.124)	(4.698)
	Adj R-sq	0.002	0.459	0.542	0.545	0.553
N(pupils)		503,746	484,647	480,036	479,972	545,904
N(schools)		2,864	2,833	2,670	2,669	3,102
KS2 score		No	Yes	Yes	Yes	Yes
NPD controls		No	No	Yes	Yes	Yes
Acorn indicators		No	No	Yes	Yes	Yes
School peers vars		No	No	No	Yes	Yes
Selective areas?		No	No	No	No	Yes

Note: ***=sig. at 1%; **=sig. at 5%; *=sig. at 10%

7.5 Pupil achievement in neighbouring schools

The failure to establish area-wide effects of GM schools on mean GCSE achievement in the previous section raises further questions as to why this might be the case. First, it is possible that no effect was found because an LA is often too large an area over

which to identify the spillover effect of an individual school. Second, even if there is no overall area-wide effect, it is possible that there is a distributional effect whereby pupils who get to attend GM schools benefit at the expense of those attending neighbouring schools. Both these additional research questions can be tentatively investigated in the dataset. However, because individual school (rather than area-wide) effects are being estimated, there is an increased risk that results are confounded by sorting on unmeasured pupil characteristics.

In this section, models of pupil achievement on schools located close to a GM school are estimated. Schools are measured in terms of the extent to which they have been affected by GM vote winning and vote taking schools, using the following school-level variables:

School SOA1: the percentage of pupils at the school who live in a lower SOA neighbourhood alongside pupils who
(i) attend GM vote taking schools;
(ii) attend vote winning schools.

School SOA2: the school's average level of pupil exposure to
(i) the percentage of pupils in the lower SOA attending GM vote taking schools;
(ii) the percentage of pupils in the lower SOA attending GM vote winning schools.

School 5km: the number of
(i) GM vote taking schools within 5 km radius of the school;
(ii) GM vote winning schools within 5 km radius of the school;
(iii) (any) secondary schools within a 5 km radius of the school.

School Near3: a binary indicator for whether one of the nearest 3 neighbouring schools is a
(i) GM vote taking school;
(ii) GM vote winning school.

Table 7.11 shows the output from these final regressions. On this occasion the preferred specification for these school-level spillover variables is the fourth column, where the school peer group is controlled for. This allows us to assess schools located close to former GM schools but who have a similar peer group to other schools, thus directly isolating the effect of the GM policy on the effectiveness of neighbouring schools.

In these specifications, a statistically significant spillover effect of GM schools on their neighbours is never found. Indeed, the point estimate values are always negative, though very small indeed. Note that the only significant coefficient in this specification identifies that being located close to any other schools (i.e. being a school in an urban area) appears to be good for school GCSE performance. This has been found by many others (e.g. Gibbons et al., 2008b), but it is not possible to say whether it is due to

Table 7.11: Effect of GM schools on neighbouring school pupil performance

		GCSE point score (best 8 grades)				
		b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)
School	GM won	8.749**	6.906***	−0.508	−0.179	1.882
SOA1:		(3.471)	(2.079)	(1.615)	(1.542)	(1.473)
	GM taken	−0.065	0.694	1.931	1.048	0.916
		(3.642)	(2.212)	(1.680)	(1.628)	(1.606)
	Adj R-sq	0.001	0.459	0.542	0.545	0.553
School	GM won	2.439	0.810	−2.843	−2.733	0.761
SOA2:		(5.534)	(3.404)	(2.813)	(2.704)	(2.538)
	GM taken	22.585***	12.129***	4.604*	2.938	2.543
		(4.784)	(2.992)	(2.477)	(2.406)	(2.325)
	Adj R-sq	0.004	0.459	0.542	0.545	0.553
School	GM won	3.206**	2.026***	0.139	−0.129	0.286
5km:		(1.295)	(0.738)	(0.595)	(0.564)	(0.504)
	GM taken	−0.567	−0.356	−0.238	−0.046	−0.194
		(1.227)	(0.729)	(0.567)	(0.544)	(0.496)
	Any school	−1.168***	−0.210*	0.291***	0.318***	0.281***
		(0.166)	(0.113)	(0.087)	(0.098)	(0.093)
	Adj R-sq	0.005	0.459	0.542	0.545	0.552
School	GM won	0.018	1.557	−1.017	−0.414	1.053
Near3:		(2.546)	(1.533)	(1.184)	(1.136)	(1.063)
	GM taken	3.893	1.996	0.169	−0.027	−0.301
		(2.439)	(1.480)	(1.114)	(1.067)	(1.032)
	Adj R-sq	0.000	0.458	0.542	0.545	0.552
N(pupils)		503,746	484,647	480,036	479,972	545,904
N(schools)		2,864	2,833	2,670	2,669	3,102
KS2 score		No	Yes	Yes	Yes	Yes
NPD controls		No	No	Yes	Yes	Yes
Acorn indicators		No	No	Yes	Yes	Yes
School peers vars		No	No	No	Yes	Yes
Selective areas?		No	No	No	No	Yes

Note: ***=sig. at 1%; **=sig. at 5%; *=sig. at 10%

competition, collaboration or another explanation.

7.6 Pupil achievement in former GM schools in 2005

This final section of analysis examines whether GM vote-winning schools directly benefit pupils who attend these schools. This is the assumption of policy-makers because all recent studies of school effectiveness using NPD show that autonomy from local authority control is associated with superior pupil performance at GCSE.

Table 7.12 shows as example of these associations by estimating five OLS regressions. Foundation schools (these are three-quarters of the former GM schools) are associated with 17 points (one grade in three of the eight best subjects or 0.17 s.d.) higher GCSE exam scores for their pupils, and this apparent outperformance of Community schools persists, even as KS2 prior attainment, other NPD indicators, geo-demographic indicators and peer controls are added. Their unexplained GCSE advantage is 2 points, or one-third of a grade in one of eight subjects (and these figures are even higher for Voluntary-Aided schools). However, the RDD approach can assess whether the apparent effectiveness of former GM schools is likely to be directly *caused* by the policy of giving schools autonomy from local authority control.

Table 7.12: OLS regression of school type effectiveness for all non-grammar schools

	GCSE point score (best 8 grades)				
	b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)
Foundation	17.490***	9.098***	7.852***	3.057***	2.127**
	(1.995)	(1.227)	(1.155)	(1.008)	(0.996)
Vol. Aided	27.794***	10.380***	9.828***	10.798***	7.641***
	(2.088)	(1.312)	(1.195)	(0.998)	(1.047)
Vol. Con.	20.389***	7.508***	7.502***	0.768	−0.685
	(3.546)	(2.313)	(2.232)	(1.754)	(1.739)
Acad. / CTC	12.356	10.163	14.188**	18.902***	16.225***
	(12.378)	(6.333)	(6.214)	(5.752)	(5.028)
Controls:					
KS2 score	No	Yes	Yes	Yes	Yes
Pupil background	No	No	Yes	Yes	Yes
Area deprivation	No	No	No	Yes	Yes
School peers	No	No	No	No	Yes
Adj R-sq	1.1%	45.0%	50.7%	53.4%	53.6%
N(pupils)	551,688	530,842	528,083	525,935	525,871
N(schools)	2,949	2,949	2,946	2,946	2,945

Note: *** = sig. at 1%; ** = sig. at 5%.

The inference that autonomy is causally related to pupil performance has been consistently claimed by government, yet it is possible that this inference is confounded by

alternative explanations. First, that Foundation schools have pupils with unobserved family circumstances that cause them to make more progress in secondary school. Second, a reverse causation argument that effective schools chose to take the vote to become GM (and therefore subsequently Foundation) schools and that these differences in effectiveness have simply persisted through to today. It is difficult to disentangle school effectiveness from unobservable pupil characteristics using the GM RDD since sorting of pupils has taken place in the past 15 years. That said, where school catchment neighbourhoods tend to have particular fixed characteristics that have persisted over the 15 years, the RDD approach does not confound these characteristics with effectiveness.

Figure 7.8 plots the average capped GCSE point score for schools (grouped to nearest 5%) who took a GM vote in the 1990s. It illustrates that there does not appear to be a clear discontinuity in 2005 exam performance between non-grammar schools that did, and did not, win the GM vote. This suggests there is likely to be no difference in effectiveness of vote winners and losers in 2005 data.

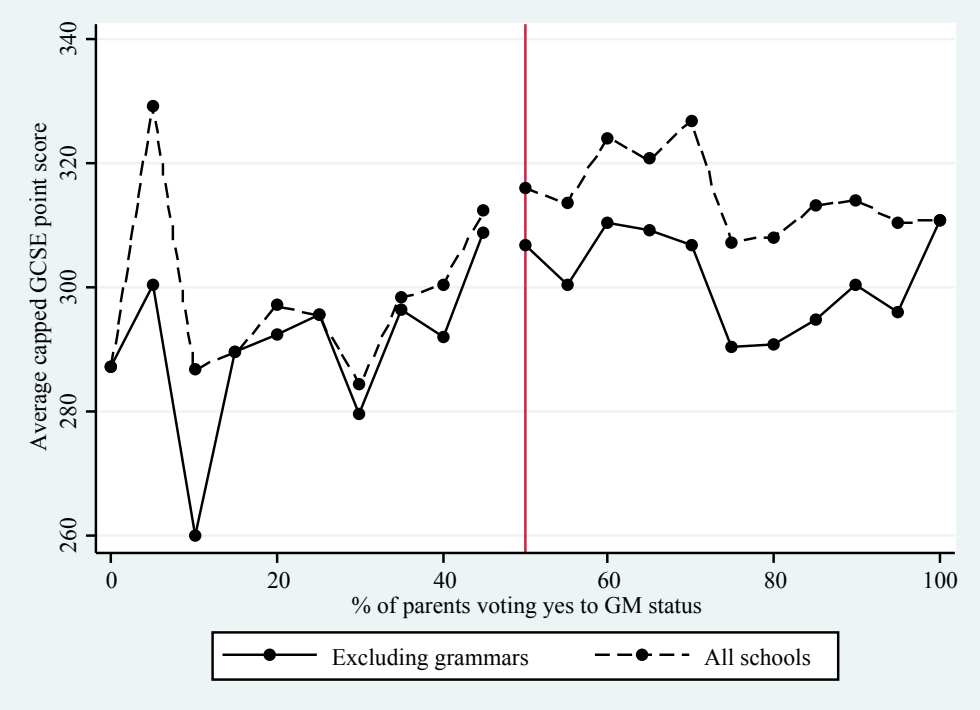

Figure 7.8: Average capped GCSE point score by GM vote

Table 7.13 shows the RDD results for a series of regressions of 2005 GCSE scores on whether the GM vote was won or not. It does this first for schools whose parental yes vote was between 35% and 65%, then for schools whose vote share was between 15% and 85%: the first set of results are less likely to have biased point estimates, but they have fewer observations so less precision. On each occasion both the simple

ITT coefficient (whether the school won the first vote) and the two-stage IV coefficient (instrumenting GM status on the first vote share result) are reported.

The first column of estimates shows that the raw differences in GCSE point scores between vote winners and losers is around 6 points (or one grade difference in one of eight subjects); this is not statistically significant. When full pupil characteristics and peer group variables are added, the gap between winners and losers is tiny – just 0.2 points – and is again not statistically significant.

Table 7.13: RDD comparing non-grammar GM vote winners and losers

		GCSE point score (best 8 grades)				
		b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)	b (s.e.)
$V = [35, 65]$	Vote winners (ITT)	5.429	1.802	1.581	0.953	0.237
		(5.013)	(2.938)	(2.746)	(2.373)	(2.297)
	Vote winners (IV)	6.468	2.149	1.879	1.131	0.282
		(5.982)	(3.507)	(3.267)	(2.819)	(2.741)
	Adj. R-sq	0.001	0.472	0.519	0.543	0.546
	N(pupils)	44,888	43,499	43,378	43,221	43,221
	N(schools)	221	221	221	221	221
$V = [15, 85]$	Vote winners (ITT)	4.991	2.124	2.534	0.495	0.316
		(3.378)	(2.064)	(1.943)	(1.656)	(1.584)
	Vote winners (IV)	5.619	2.392	2.852	0.558	0.356
		(3.802)	(2.324)	(2.187)	(1.867)	(1.787)
	Adj. R-sq	0.001	0.455	0.504	0.530	0.534
	N(pupils)	118,919	114,922	114,597	114,141	114,141
	N(schools)	593	593	593	593	593
Controls:	KS2 score	No	Yes	Yes	Yes	Yes
	NPD controls	No	No	Yes	Yes	Yes
	Acorn indicators	No	No	No	Yes	Yes
	School peers vars	No	No	No	No	Yes

The OLS coefficient on Foundation status from a regression with a full set of controls is ten times larger than the (insignificant) RDD coefficient on being a former GM school *in exactly the same dataset*. This suggests that there is no evidence that this policy of school autonomy produced more effective secondary schools. It is possible that Foundation schools are effective schools, but so are schools that remained in LEA control because they lost the GM vote. If so, this tells us more about the type of schools that elected to hold a parental vote than the causal effect of a policy, *per se*. More likely, there is no genuine difference in the effectiveness of LEA controlled and autonomous schools, with apparent effectiveness of Foundation schools attributable to unmeasured characteristics of pupils in these schools today.

7.7 Discussion

Education policy reforms are usually made without a randomised element so that quantitative researchers struggle to find a good counterfactual for the policy not having taken place. This undermines our ability to draw conclusions about whether the policy works. In this respect the GM policy was unusual: the desire to legitimise the policy through a vote of parents (unintentionally) created a discontinuity whereby apparently similar schools just did, or did not, receive the treatment. This has enabled us measure the direct and the spillover effects of the policy both in the short-term and in the long-run.

Many commentators on education policy regard the Grant-Maintained schools policy as a historical experiment in autonomy: it was abolished ten years ago and several of the powers these schools once had have been removed (or given to all Community schools). This chapter is motivated by the belief that the effects of this policy have persisted and can been measured in current secondary schooling data. Former GM schools, with Foundation (or VA) status, are able to retain any changes made to their admissions policies in the 1990s and they can still administer the process of pupil admissions. This policy decision is reflected in the evidence that former GM schools do have pupil intakes slightly higher in ability and income status than schools who lost their GM vote and so remained as Community schools.

Foundation schools also still differ from Community schools in the amount of control they have over staffing and spending decisions. The assumption that this autonomy from local authorities is desirable is implicit in many strands of recent government reform. This chapter suggests that any new policies that give schools autonomy without other major institutional changes are unlikely to lead to sustained improvements in pupil exam performance since former GM schools perform no better than vote-losing schools, once pupil background is taken into account.

More specifically, taken together with Clark (2007), the studies suggest that the introduction of GM status may have led to quite substantial improvements in pupil achievement in the short-term, but these have not persisted to today. GM status emerged alongside the first publication of school league tables in 1992, and these schools may have responded more quickly than others to the new incentives to maximise the proportion of pupils scoring 5+A*-C at GCSE (West and Pennell, 2000). There might be technical reasons for our differences in short-run versus long-run findings (most notably differences in GCSE outcome measure and slightly different samples of schools), or they might be explained by the slightly reduced levels of autonomy Foundation schools enjoy. However, it seems perfectly possible that our different findings reflect real substantive differences in how policy implementation affects schools as time pro-

gresses.

Changes in school governance may enable substantial school improvements in the short-term as it encourages a fresh start for management and teachers with renewed focus on pupil achievement. However, in the long-term, this higher level of effort directed at exam outcomes (relative to other schools) starts to decline. Case studies looking at leadership, management and teaching in schools in schools over the policy cycle would help us understand why this might happen.

This chapter also fails to lend support to the assumption underlying market reforms that structures to encourage schools to compete for pupils are a route to improving standards. It is likely that secondary schools adjacent to former GM schools do still perceive them to be a competitive threat since Foundation schools recruit pupils from a much wider catchment area than their direct neighbourhood. However, consistent with the short-run findings of Clark, this apparent competition for pupils (if there is any) does not appear to translate into effort on the part of neighbouring schools to increase exam performance.

Chapter 8

Spillover effects of religious schools

Religious schools, maintained by the State, are an important feature of the English schooling system, educating 15 per cent of secondary-aged children. They are a legacy of the 1902 settlement between the Government and the Church of England (CofE) and Roman Catholic (RC) churches, as the principal providers of 19th century schooling. Despite a steep decline in church attendance across all denominations and limited support for the principle of state-funded religious schooling in attitude surveys, there continues to be relatively high demand for places at religious secondary schools, with institutional rigidities constraining the supply of places.[1]

Religious secondary schools provide a genuine and enhanced opportunity for some parents to actively choose between faith and non-faith schools (without the cost of moving house) since they usually give priority in admissions based on religious affiliation of parents rather than proximity of home to school. This was demonstrated in Chapter 4, which showed that far fewer pupils attend their nearest school in areas where there are many religious secondary schools. This choice means that, in a system with spare capacity, religious schools do appear to present an enhanced competitive threat to neighbouring schools, who may respond by exerting effort in some way to attract local families to their school. However, for choice and competition to lead to higher achievement in schools, the incentives for head teachers and teachers need to be aligned such that all involved respond to competition by making efforts to 'raise standards', however that may be defined. These incentives structures may be weak where there

[1]For example, in 2006, 62% of respondents agreed with this question in a Populus poll: '*Faith schools are divisive because they prevent children from different religious backgrounds from getting to know and understand each other*'. In 2005, 64% of respondents agreed with this statement in an ICM poll: '*Schools should be for everyone regardless of religion and the government should not be funding faith schools of any kind*'.

are few changes in capacity in the local educational market so that school survival is guaranteed regardless of quality. They will also be weak if parents judge schools based on league table position of overall achievement, rather than measures that take account of the quality of the schools intake. Under these circumstances, schools can best ensure they survive and prosper simply by concentrating their effort on securing an advantaged intake through the application of certain admissions policies and procedures (i.e. by cream-skimming more able or easier to teach pupils). This suggests that a quasi-market for school places risks becoming quite stratied.

This chapter measures the extent to which the presence of religious schools 'spills' over to the educational experiences of pupils who attend neighbouring schools, whether through school effort induced by competition, changes in school sizes induced by choice or changes in peer groups induced by sorting. The estimation strategy combines, and improves, methods proposed by Card et al. (2007) and Hoxby (1994) to identify these effects without the confounding influences of unmeasured pupil characteristics. Pupil growth in achievement from KS2 to KS4 is modelled as a function of share of pupils in the area who attend religious schools, using pupil fixed effects and a wide range of pupil and area controls. In addition, instrumental variable methods are employed to avoid confounding the causal effect of religious schools on GCSE outcomes with direct effects of religious families on educational success, or with growth in religious schools caused by ineffective local Community schools. The historic characteristics of an area – early 20^{th} century levels of religious affiliation – are used to predict the modern-day supply of Catholic school places, thus isolating a source of variation that is exogenous to current demand for religious school places.

An important part of the 'whole area' effect of religious schools will be the relative effectiveness of religious schools compared to local authority schools. In England, Voluntary-Aided secondary schools outperform other comprehensive schools, though by relatively little (just under one grade in one subject at GCSE) once all observable characteristics of the pupil intake are accounted for (see Table 7.12 in Chapter 7). This estimate is roughly consistent with Schagen and Schagen (2005) who agree that the higher performance of religious schools at GCSE and KS3 is almost entirely attributable to the nature of their intake. Gibbons and Silva (2006) employ an innovative method to deal with the confounding influence of sorting based on unobservable characteristics into religious primary schools. By comparing families who switch between religious primary and non-religious secondaries with families who transfer from non-religious primary to religious secondary school, they are able to show that the first group make no greater educational progress during primary school than the latter group.

There also exists an extensive literature on the effectiveness of US Catholic schools,

operating in the private sector, and the methodological controversies that arise from this literature are relevant to this chapter. Early studies of schooling in the US by Coleman et al. (1982) and Coleman and Hoffer (1987) suggested that Catholic schools have positive effects on test scores and high school graduation rates, which were said to have come from their academic structure and culture, their internal community, their devolved governance, and the inspirational ideology (Bryk et al., 1993). However, early studies did not account for the selection bias that results from factors affecting the probability of attending a Catholic school (such as religious affiliation and income) also directly entering the education production function. There have been numerous attempts to account for this selection bias through the use of instruments including Catholic faith, distance to nearest Catholic school and the interaction between these two (e.g. Evans and Schwab, 1995; Grogger and Neal, 2000; Jepsen, 2003; Neal, 1997; Sander, 1996). However, Altonji et al. (2005) showed that most of the instruments used to predict attendance were not orthogonal to pupil achievement, even in value-added achievement models. This means that effect sizes are likely to be biased upwards because Catholic schools tend to select children from high-religiosity families and activities within these homes are known to produce better educational outcomes, regardless of school attended (e.g. Bankston and Zhou., 2002; Regnerus and Elder, 2003; Glanville et al., 2008).

In recognition that no suitable instruments to predict Catholic school attendance have been found, Nguyen et al. (2006) use propensity score matching to compare the educational attainment of observably-identical children at Catholic and non-Catholic schools, finding that US Catholic schools raise test scores (but only for males in maths), graduation rates and likelihood of enrolment in 4-year college. However, matching relies on the condition independence assumption, which states that the outcome of non-participants is assumed to be independent of participation, conditional on a vector of observable characteristics (Rosenbaum and Rubin, 1983). In other words, they again rely on the assumption that Catholic families have no unobservable characteristics that raise child's achievement or progress, independent of school attended. Thus, it is not clear why a matching approach, even using a value-added specification, would overcome the problem of sample selection bias.

There are two papers estimating whole area effects of Catholic schools in the US and Canada that are directly comparable to this analysis. Hoxby (1994) uses the National Longitudinal Survey of Youth to estimate the effect of the presence of Catholic schools (which constitute over 80 per cent of the US private schooling market) on area-wide achievement. The supply of Catholic schools is instrumented using the current size of the Catholic population in the area. She finds that a 10 percentage point increase in Catholic school enrollment produces 0.9 additional years worth of educational achievement, on average, and 6% higher wages for pupils graduating from

high schools in the area. However, her choice of instrument is controversial, even though she is able to control for the religious affiliation of the family, because it requires us to assume that Catholic families who live in predominantly Catholic areas are no different from Catholic families who live in areas with few other Catholic families.

Canada's dual system of schooling is similar to that of England because Catholic schools are entirely state-funded on the same terms as secular schools. Card et al. (2007) estimate the effect of Catholic school enrolment share on test score gains by pupils, using a school fixed effects specification of test score growth between grades 3 and 6. They find small positive effects from competition in the area of Ontario, but have no means to isolate an exogenous source of variation in the supply of Catholic school places. This means that their estimation strategy relies on comparisons between areas with different fractions of Catholic families to identify cross-system effects. This again raises the concern expressed by Altonji et al. (2005) that the local religious composition of the population enters the education production function via some other mechanism. However, although they have no measure of the religion of the child's family, the identification assumptions would seem less onerous than Hoxby's because they use a value-added specification combined with school and cohort dummies.

The remainder of this chapter begins by documenting the history of church attendance and religious schooling in England - a history that is particularly important given the use of historical instruments in the estimation strategy. Next, theoretical models of competition presented in Chapter 6 are applied to describe how religious schools might impact on whole area GCSE achievement. The estimation strategy is then described and assumptions necessary for successful identification are discussed. The use of a pupil fixed effects specification is a direct improvement on the school fixed effects approach used by Card et al. (2007). The use of historical instruments, combined with controls that include modern-day religious populations, substantially reduce the identification assumptions required by the Hoxby (1994) IV approach. Next, the data sources are summarised (more detailed can be found in Chapter 2). Finally, the results are estimated and discussed.

8.1 The supply of religious secondary school places

This section describes the factors contributing to growth in the provision of CofE and RC secondary schooling in England. Changes in the geographical distribution of Anglican and Catholic communities across England are important determinants of both historic secondary school supply and modern-day demand for faith school places. Catholic schools are given special attention since they comprise about two-thirds of all religious secondary schools today, and their effects are estimated separately in the

results sections of this chapter.

8.1.1 Growth and decline of church attendance

The first census of places of worship took place on March 30th 1851, giving a geographically detailed description of the relative sizes of Anglican, Catholic and non-conformist communities. The census reported that (no more than) 58% of England's population attended church on the Sunday in question, with over half of those in Church of England congregations. Non-conformists groups made up the majority of the rest of church-goers, with one-in-five church attendees at a Methodist service, just under one-in-ten at Baptist services and one-in-ten at an independent church service.

According to this census, levels of formal worship in Roman Catholic churches in England were very low at around 2% of the English population. There were large geographical disparities, with Catholic communities concentrated around the north of England (e.g. Lancashire, Northumberland, Staffordshire, Durham) and the Midlands (Warwickshire, Leicestershire). Other areas such as Rutland, Huntingdonshire and counties in the East of England had almost no Catholic church-goers. These Catholic communities were, for the most part, recent immigrants to England. The size of the community had grown from around 100,000 in 1800 to 250,000 by 1840 as a result of Catholic emancipation and Irish immigration. This increased further following the Irish potato famine in 1845, but the 1851 census failed to capture much of this mass immigration.

Survey data indicates that the size of the Catholic population in England grew substantially, with around 18% of the population identifying themselves as Catholic by 1961. Figure 8.1 also shows that Catholics became more geographically dispersed across ancient counties (large geographical sub-divisions of England used prior to 1965) from the mid-19th to the mid-20th century. However, some caution is needed regarding the quality of this survey data because no church censuses took place over this time, meaning we have no measure of the strength of religious affiliation of these people. This is important because this historical Catholic data predicts the modern-day supply of Catholic schools, but the actual size of the church-going population would be a better predictor of the wealth of the local church, and therefore its ability to finance new Catholic schools.

The churches resumed holding censuses of church attendance in 1979, 1989, 1998 and 2005, giving some indicator of modern-day demand for religious schooling. These censuses documented the very rapid fall in church-going, with overall church attendance (on the day of the census) declining from 11.7% in 1979 to 6.3% in 2005. Anglican attendance fell from 3.6% to 1.7% and Catholic attendance fell from 4.3% to 1.7% of

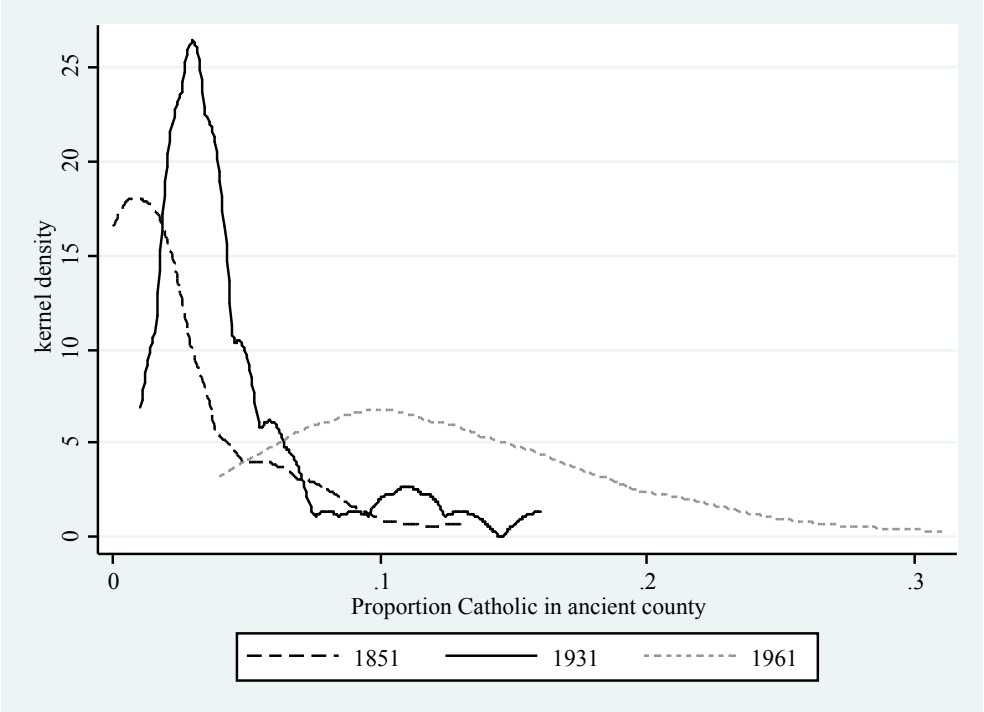

Figure 8.1: Changes in dispersion of Catholic population across England

the population for the same period.

These surveys of church attendance also reveal the current geographical distribution of church-goers. Catholics are still predominately located in London and in the northern urban areas of Lancashire, Merseyside, Cleveland and Durham. However, their concentration in these areas is less pronounced than in 1851. By contrast, the Church of England has a relatively low presence in the industrial areas of London, South and West Yorkshire, Merseyside, Tyne and Wear, Cleveland and Greater Manchester.

There exists an alternative measure of religious allegiance that provides a different of indication of likely demand for religious schooling. In the 2001 Census of Population, a question about religion was asked, namely: '*What is your religion?*'. The options given were none; Christian; Buddhist; Hindu; Jewish; Muslim; Sikh; any other religion. Altogether, 71% of the population reported being Christian, with 15% reporting no religion, 8% not stating a religion, 3% reporting that they were Muslim, 1% each Hindu, Jewish and Sikh. Unfortunately no details on denominational allegiance are revealed in this population, but it does suggest that a much larger group of families might be attracted by a religious education for their child than is representative of church-going levels.

8.1.2 The supply of religious secondary schools

The churches, and in particular the established Church of England and Roman Catholic churches, have played a highly significant and enduring role in the provision of education in England. Before 1870, most schooling took place without state involvement. Most early voluntary schools had religious (both Anglican and non-conformist) underpinnings and purpose, although the influence of secular philosophers was also a feature of the late eighteenth and early nineteenth century education system (Johnson, 2006). The contribution of the churches was not always motivated by simple social altruism, but instead was at times seen as a means of promoting and retaining particular denominational allegiance (Skinner, 2002). The Church of England (CofE) had long been involved in the education of the social elite in 'public' (i.e. fee-paying private) schools, but in 1811 it established the National Society for the Education of the Poor as a strategic response to competition with the nonconformist churches, especially the Methodists, for the loyalties of the English working class (Johnson, 2006). The initiative to provide an Anglican education for every parish in the country was contested by dissenting churches, liberal Anglicans and some Jews, who believed that schools should not be specific to one denomination. So, from 1814 onwards, the British Foreign School Society for the Education of the Labouring and Manufacturing Classes of Society of Every Religious Persuasion began set up elementary schools, with a non-denominational Christian curriculum. A tiny minority, represented by the Central Society of Education, favoured schools with no reference at all to religion (Gardner et al., 2005).

The setting up of Roman Catholic schools started in 1847, following mass Irish emigration into the industrial conurbations of London, Liverpool, Lancashire, Birmingham, Manchester, Newcastle and the North-East. This network of Catholic schools were intended to be '*a cultural and faith bastion against the potentially polluting effects of hegemonic Protestantism and secular rationalism*', aiming to provide a place in a Catholic school for every Catholic child (Grace, 2002, page 8). The distinct Catholic educational mission of this time was to the poor; it is notable that the precursor to the Catholic Education Service was called the Catholic Poor School Committee (Grace, 2002). However, there were also a small number of Catholic independent and grammar schools established to provide leadership for the sons and daughters of the Catholic upper and middle classes.

Although the British government was initially reluctant to be involved in education, it financially supported the voluntary sector from a very early stage. Grants were increased rapidly between the years 1833 and 1857, and so to administer the large sums a Department of Education was set up (Johnson, 2006). State involvement in the provision of schools began under the 1870 Education Act, which saw the establishment

of the 'Dual System' of elementary schooling provision. The intention was that the churches would provide the majority of school places and school boards would be set up to fill gaps in provision (Barber, 1994). The election of members to the school boards became a partisan battleground for Nonconformists against the Established Church as they attempted to transform the political and social framework of the nation in the hopes of obtaining their concept of complete religious equality (Richards, 1970). The response of the National Society (Anglicans) and Catholics to the threat of board schools displacing the need for church schools was to move at speed to raise money to build new schools in the grace year before board schools could be opened. In the 15 years following the 1870 Education Act, CofE schools increased from 6,382 to 11,864 and RC schools from 350 to 892 (Gardner et al., 2005). This meant that by the turn of the century, nearly three-quarters of the 20,000 elementary schools were under religious or voluntary bodies (Ministry of Education, 1950). However, this rapid growth in voluntary schools provision produced financial difficulties within both the Anglican and Catholic sectors that would continue into the 20th century (Gardner et al., 2005).

Financial difficulties forced the churches relinquished some of the control over their schools as part of the fiercely contested 1902 Education Act. Their schools were now overseen by local authorities, but in return they were relieved of some of the burden of paying for them (e.g. teacher salaries and wear and tear on buildings). Church schools would continue to receive a capital maintenance grant of up to 50 per cent, but once the system was in place there would be no new building money (Gardner et al., 2005). This Act also gave power to local authorities to set up secondary schools for the first time. There was consistent opposition to state subsidies for church schools throughout the debating of the 1902 Act, with many arguing that the voluntary sector should be abolished or allowed to wither away (Johnson, 2006). Nonconformists, who had few schools of their own because most were ideologically opposed to a sectarian education system, were incensed at the extent of state subsidy to the Anglican and Catholic schools, particularly as in many rural areas there was no state-funded alternative to the local CofE school (Barber, 1994).

State support for church schools was not sufficient to remove financial problems. At the outbreak of WWII, churches provided elementary education for around a third of children. However, the church schools were in a state of disrepair – in 1939 the government's black list of schools with defective premises had 735 schools on it, of which 72 per cent were church schools. Furthermore, only 28 per cent of church elementary schools had implemented the Hadow reforms to partition pre- and post-11 education (Board of Education, 1941). The church schools' lack of money and inefficiency due to their small average size was widely seen as a hindrance to educational reorganisation, which caused most of the Liberal and Labour political élite to advocate outright

abolition of the sector. This reluctance by the churches to support any educational reform led many to claim that they were putting their own self-interest before the national interest (Barber, 1994). The Roman Catholics faced the most severe financial problems and were the most difficult for Butler to negotiate the 1944 Education Act with, whereas many in the Church of England took a more pragmatic approach to the provision of education (Barber, 1994).

Following the 1944 Education Act a system of universal, free and compulsory education was set up. Section 11 of the Act required that every LEA should prepare a plan for primary and secondary schools, in consultation with diocesan educational authorities. Church schools were offered a chance to become 'Voluntary Controlled' (VC), meaning that the State would take day-to-day control of the school. If the church wished to retain control of their school via 'Voluntary Aided' (VA) status, they were required to make a 50% capital contribution to the school (this figure was reduced to 25% in 1959 and is now 10%). VA status retained significant church involvement and control over schools. They had majority representation on the governing body, giving them control over the employment of staff, buildings and repairs, and school admissions (Gay and Greenough, 2000).

All Roman Catholic schools opted for voluntary aided status, but the Church of England gave no clear national guidance as to whether schools should pursue VC or VA status, and both R.A. Butler (then Secretary of State for Education) and the Archbishop of Canterbury expected few schools to opt for VA status (Barber, 1994). In a survey of key diocesan decision-makers at the time, Kelly (1978) found some dioceses making strong statements in support of aided status (Bath and Wells, Blackburn, Carlisle, Chester, London, Portsmouth, Southwark, Winchester and York) and others equally strongly supporting controlled status (Bristol, St Edmundsbury and Ipswich, Sheffield). Other diocese either gave a less clear lead, or allowed decisions to be made entirely locally, and this de-centralised decision-making led to the geographical diversity in the provision of Anglican schooling that continues today.

Most church-run elementary schools converted to primary schools, for two reasons. First, neither the Anglicans nor the Catholics had the resources to provide facilities suitable for secondary education (Ministry of Education, 1950). Second, in some areas the will to build new secondary schools was seen to be lacking: it was difficult to get local Parish financial support for secondary schools since they recruited from a very large geographical area (Gay and Greenough, 2000; Ministry of Education, 1950).

The 1959 Education Act allowed for the possibility of new church secondary schools via increased government financial support for capital costs through both loans and larger grants (Church of England Board of Education School Committee, 1972). The Catholics (much more than the Anglicans) made extensive use of government loans to

embark on a school building programme in the 1950s and 1960s, resulting in a growth in the number of secondary aged pupils at RC schools between 1950 and 2000 from 50,000 to 309,000. So, they were at partly able to keep up with the population diffusion of the Catholic population from its earlier nineteenth-century concentrations into the new towns and suburbs of the twentieth century (Gay and Greenough, 2000).

Expansion in the Anglican secondary school sector was modest, with an increase of secondary aged pupils from 64,000 in 1950 to 150,000 in 2000. This was due to financial difficulties that continued in the decades following the 1944 Education Act. By the early 1970s, inflation was making it difficult for them to meet their share of the capital costs of VA schools. However, within the Anglican community, VC status was now seen as less acceptable than it was in 1944 (partly because the amount of time devoted to religious instruction had fallen to 2 periods a week), so it was preferred that VA schools were allowed to close, with capital reinvested elsewhere (Church of England Board of Education School Committee, 1972). However, although national recommendations were made to Anglican schools there was never any attempt to harmonise provision across the country (Church of England Board of Education School Committee, 1972).

The final decades of the 20^{th} century saw very little change in the supply of faith schools in England, but there is the prospect of a new rapid expansion in the sector. The White Paper '*Schools Achieving Success*' called for the expansion of faith-based education, with the most generous financial conditions ever offered by a government to faith communities who were now required to make just a 10 per cent capital contribution to new schools (Department for Education and Skills, 2001) . In return for this generosity, the White Paper sought to achieve greater mixing between religious (and non-religious) communities within schools by suggesting that 20 per cent of a faith school's places should be set aside for pupils of other faiths, or none (Gardner et al., 2005).

8.1.3 The distribution of religious schools today

Table 8.1 shows the current religious denominations of secondary schools in England. Around two-thirds of these are Roman Catholic. Whilst overall church attendance is comparatively low in England, attendance at religious schools is significantly higher. Thus, in England as a whole, although only 6% of the population attended church one Sunday in May 2005, over twice as many pupils (15%) attended a religious secondary school in 2005 (see Table 8.2). RC schools are over represented, relative to church-going populations in almost all areas, but particularly in Lancashire, Durham, greater London, Birmingham. Figure 8.2 shows there remains wide variation in religious school supply across England.

Table 8.1: Religious denomination of secondary schools

	N	%
None	2,577	82.9
Church of England	151	4.9
Roman Catholic	338	10.9
Jewish	7	0.2
Muslim	2	0.1
Seventh Day Adventist	1	0.0
Church of England/Roman Catholic	5	0.2
Roman Catholic/Church of England	2	0.1
Christian	23	0.7
Church of England/Christian	1	0.0
Sikh	1	0.0
Total	3,108	100.0

Table 8.2: Proportion of pupils in religious schools by ancient county

	Mean	Std. Dev.	Min.	Max.	N
Proportion in religious schools	15.0%	7.4%	0.0%	30.9%	39
Proportion in CofE schools	4.5%	3.2%	0.0%	16.6%	39
Proportion in RC schools	9.5%	6.1%	0.0%	22.6%	39

Weighted by ancient county size

This reflects the historically high levels of Catholic affiliation in these areas, which combined with state subsidies has meant that supply has not fallen with a decline in church-going. In any case, many of these schools remain very oversubscribed, not necessarily because parents are attracted to the religious education, but because they recruit a more advantaged intake than other schools in the area. Survey evidence suggests that the relationship between levels of Catholicism and the supply of RC school places is not now strong. Grace (2002) reports that only 12 out of 60 Catholic schools surveyed in London, Birmingham and Liverpool had 100 per cent Catholic student enrolment. At 5 of the 60 schools, Catholic students constituted only 50 per cent of the enrolment.

8.2 Competition from religious schools

Chapter 4 estimated that if one quarter of pupils in a LA are educated in a religious school, this raises the number of pupils who attend a non-proximity school by 15 percentage points. However, though religious and non-religious schools do appear to actively recruit pupils from the same neighbourhoods, this empirical observation is not sufficient to assert that they actively compete with each other for pupils. Following the

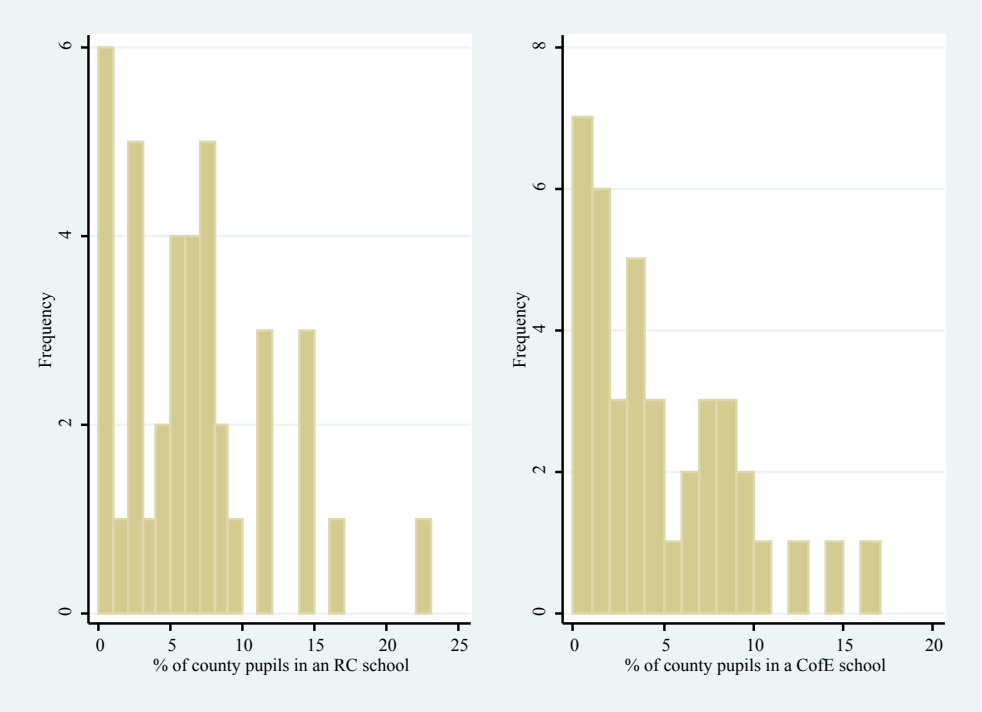

Figure 8.2: Proportion of pupils in RC and CofE schools across ancient counties

economic models of school competition discussed in Chapter 6, the level of competition a non-religious secondary school faces from the presence of a nearby religious school can be summarised by the Bayer and McMillan (2005) concept of *residual elasticity of demand* – the changes in a school's demand in response to a change in only that school's quality. If the residual elasticity of demand at a non-religious school is raised by the presence of local religious schools, it must be true that many of the school's potential parents have a feasible choice set that also includes religious schools. Recalling that the definition of feasible choice set used in this thesis requires the schools to be willing to accept the child if they apply, these families would usually need to be capable of demonstrating religious adherence through church attendance. It is difficult to measure the proportion of families who actively consider both religious and non-religious schools, and thus contribute to the competitive threat.

It is a necessary, but not sufficient, competition condition that families in the area have both religious and non-religious schools in their feasible choice set. A second requirement for a school's residual elasticity of demand to be raised by a religious school is that parents for whom both schools are in the feasible choice set would be responsive to a change in the school's perceived quality, whether academic or otherwise. This would be the case where parents value academic results or reputational information very highly, relative to other characteristics such as distance to home or religious ethos. Where religious ethos of school dominates the household utility function, there is

little opportunity for non-religious schools to attract these parents by increasing school quality. Schools' residual elasticity of demand will be higher where schools are more closely matched in terms of league table position. One reason why religious schools might not increase the level of competition in an area is that they are associated with greater ability stratification of schools (see Chapter 4), so all schools' residual elasticity of demand is consequently lowered.

Finally, competition between religious and non-religious schools will be stronger where the parents who are actively considering both schools have children who are perceived as desirable to teach. Of course, for a school with spare capacity, all pupils are desirable in the sense that they bring additional resources to the school, but many schools are over-subscribed so that they are only incentivised to compete over pupils who are likely to improve their league table position. Data from Chapter 5 showing that the higher ability and SES children from neighbourhoods are more likely to attend religious schools suggests that they do threaten the peer composition of non-religious schools.

Of course, as with all studies of school competition, raising effort focused on exam scores is not the only route to attempting to raise *perceived* school quality, especially where parents principally use a school's league table position as an indicator of quality. Alternative responses such as recruiting higher ability pupils or increasing marketing efforts would not manifest themselves in improved area-wide GCSE outcomes, even if competition were genuinely significant.

8.2.1 Secondary school transitions

A competition effect arising from the presence of religious schools will only be evident where there are religious families who are willing to switch between religious and secular systems to access higher quality schooling. Therefore, one way to seek confirmatory evidence on whether religious schools are likely to raise the competition is to look at the data on pupil transitions from primary to secondary school, and also between secondary schools for the minority of pupils who make these transitions.

Primary to secondary transitions

The major point of competition between schools is for pupils as they make their transition from primary to secondary school at age 11. Chapter 3 showed that at this transition point parents rely on league tables and social networks to form their opinions about school quality. If religious and non-religious secondary schools do genuinely compete for pupils, we should expect to see some transition between these two sectors

as parents rely on assessments of secondary school quality, as well as preferences for religious ethos to influence their choices.

Table 8.3 shows that transitions between the CofE and non-religious sectors are very high in both directions. By contrast, a large majority of children who are in RC secondary schools also attended RC primary schools. The reverse is also true. This may partly be due to the similar size of the RC primary and secondary sectors; to the use of feeder school admissions criteria in the RC sector, which act as a barrier to entry from secular primary schools; or may reflect a stronger preference for a religious education amongst Catholics than Anglicans. The implication of this is that the competitive threat posed by the presence of Catholic schools may be lower than for other religious schools.

Table 8.3: Primary to secondary school transitions (%)

	Secondary school					
	None	Church of England	Roman Catholic	Other Christian	Other religious	Total
Primary school:						
None	66.82	2.71	1.64	0.48	0.05	71.70
Church of England	15.70	1.64	0.51	0.21	0.00	18.06
Roman Catholic	2.01	0.13	7.33	0.10	0.00	9.56
Other Christian	0.43	0.07	0.02	0.01	0.00	0.54
Other religious	0.03	0.00	0.00	0.00	0.11	0.14
Total	84.99	4.55	9.50	0.80	0.16	100.00

Transition data is available for 96.5% of the sample

Secondary to secondary transitions

Transitions between secondary schools are relatively rare in the English system. 4.7 per cent of pupils in the sample moved schools between years 8 and 11 (there is no data collection available for these pupils in year 7), but most of these moves would seem to be motivated by family house moves. Just 1.3 per cent of the pupils in the sample move schools while retaining the same postcode in years 8 and 11. Where these moves take place that are likely to be motivated by school quality, they should be better informed decisions than those made at the primary-secondary transfer since parents would now have improved knowledge on school quality and factors influencing their child's well-being.

Table 8.4 shows that for pupils who move secondary school, there are substantial numbers of moves between the secular and religious sectors between the ages of 12 and 15. At the county level, school moves between year 8 and 11 appear to be negatively

Table 8.4: Secondary school moves between age 12 and 15 (%)

	Age 15/16					
	None	Church of England	Roman Catholic	Other Christian	Other religious	Total
Age 12/13:						
None	79.79	3.39	3.16	0.80	0.06	87.20
Church of England	3.09	0.24	0.15	0.03	0.00	3.52
Roman Catholic	6.11	0.31	2.17	0.05	0.00	8.64
Other Christian	0.53	0.02	0.02	0.01	0.00	0.59
Other religious	0.03	0.00	0.00	0.00	0.02	0.05
Total	89.55	3.96	5.51	0.90	0.08	100.00

Transition data is available for 96.5% of the sample

correlated with proportion in RC (but not CofE) schools, but this is likely due to their lack of presence in rural areas where children are forced by geography to move schools when their family moves house. For the sub-set of children who only move school and not house, these moves are positively correlated with the proportion in faith schools in the area ($\rho = 0.26$). Over 80 per cent of these moves are between two secular schools, with very few moves between two religious schools. There are slightly fewer moves from a secular school to a religious school (4 per cent to CofE and 5 per cent to RC) than there are from a religious school to a secular school (3 per cent from CofE and 9 per cent from RC).

8.3 Method

This section describes the pupil fixed-effects estimation strategy to identify area-wide effects of religious schools. Following the notation in Chapter 6, this chapter assumes the education production function for child i in school j in area k at time t can be represented as:

$$Y_{ijk.t} = \gamma_0 + \gamma_{1t}\theta_{jk} + \gamma_{2t}C_{jk} + \gamma_{3t}R_{jk} + \gamma_{4t}U_{jk} + \gamma_{5t}x_{ijk} + \epsilon_{ijk.t}$$

That is, test achievement for child i at time t, $Y_{ijk.t}$, is a function of the school cohort's peer group, θ_{jk}, the school's effort induced by competitive pressure, C_{jk}, any effectiveness associated with the school having a religious denomination, R_{jk}, all other school effort not related to competition, U_{jk}, and the measured time invariant, x_{ijk}, and unmeasured, $\epsilon_{ijk.t}$, characteristics of the child. All these time-invariant characteristics of the school and pupils have the potential to affect both the level and rate of growth of pupil test scores, and are thus modelled as time variant.

The presence of religious schools has the potential to:

1. change efficiency via the presence of a school that may be differentially effective as a result of its religious status, R_{jk}, and via competition effects, as measured by C_{jk}, for any schools;
2. change the relative sizes of schools, which affects mean area-wide achievement if differentially effective schools grow and shrink;
3. produce student sorting effects that change the peer groups, as measured by θ_{jk}, at each school.

It is difficult to separately measure these three effects since many characteristics of the child that affect their educational achievement are unaccounted for in NPD. Instead, the overall aggregate effect of the presence of religious schools is the principal parameter of interest in this chapter. By comparing the area-wide effect of religious schools (scenario 1), compared to a counterfactual where there are no religious schools (scenario 0), assuming no sorting of pupils in or out of the area, the mean effect on pupil achievement in area k is:[2]

$$\frac{1}{n}\sum_{i=1}^{n}(GCSE_{ik}^{1} - GCSE_{ik}^{0}) = \frac{\gamma_1}{n}\sum_{i=1}^{n}(\theta_{ik}^{1} - \theta_{ik}^{0}) + \frac{\gamma_2}{n}\sum_{i=1}^{n}(C_{ik}^{1} - C_{ik}^{0}) + \frac{\gamma_3}{n}\sum_{i=1}^{n}R_{ik}^{1}$$

This chapter identifies these spillover effects of religious schools using a model of pupil test score growth from Key Stage 2 (KS2 – secondary school entry at age 11) to KS4 (compulsory school leaving at age 16). It identifies whether steeper test score trajectories are associated with a greater presence of religious schools in the area. Pupil and area variables that might influence the growth trajectory are used as controls, including denominational church-going and other religious affiliation levels in the county. The key identifying assumption is that variation in the supply of religious schools (not explained by variation in the current size of the religious population) has no direct effect on test score gains of the students in the county, except via the competition mechanisms described above. This approach is similar to that used by Card et al. (2007) in their paper on Catholic schools and competition in Canada. However, the advantages of the English data are that linked repeated pupil measures of achievement allow pupil fixed-effects to be used, whereas their paper relies on school fixed-effects. Thus, more variation in test scores is explained by the fixed effects and the potential for unexplained family characteristics biasing the results is reduced.

In the basic pupil fixed-effects model, the achievement of the child, Y_{ijk}, at KS2 and KS4, is modelled as a function of the area proportion of pupils in religious schools (interacted with time dummies), $\%religsch_k$, the observed characteristics of the area

[2]School subscripts are removed since pupil i may attend a different school j in scenarios 0 and 1.

(W_k) and pupil (X_{ijk}) (all interacted with time dummies), a set of pupil fixed effects, ξ_{ijk}, to capture the time-invariant characteristics of the pupil, and a time-variant residual, $\epsilon_{ijk.t}$. $\%religsch_k$ is the proxy for the level of cross-system competitive pressure due to religious schools.

$$\begin{aligned} Y_{ijk.t} = \%religsch_k.KS4.\delta_1 + W_k.KS4.\delta_6 \\ + X_{ijk}.KS4.\delta_7 + KS4.\delta_8 + \xi_{ijk} + \epsilon_{ijk.t} \end{aligned} \tag{8.1}$$

In the second pupil fixed-effects specification, an identifier for whether the school is religious ($relig_{jk}$) and school-level variables measuring the school peer group and other characteristics (Z_{jk}) are added in an attempt to explain part of the pupil test score growth trajectory. This should have the effect of removing the direct effect of attending a religious school or a school with a particular peer group from the model, and thus the coefficient on $\%religsch_k$ reflects a more pure competition effect. However, we have no reason to believe that the actual coefficients on these school variables are unbiased, since they also capture average test score growth at the school that is due to systematic pupil sorting on unmeasured characteristics in the area. This might be an appropriate specification for isolating a competition effect if any unmeasured pupil characteristics at religious schools arise directly from the religiosity of their families. However, if the unobserved characteristics arise from a social selection effect whereby the more affluent pupils on a street attend a religious school and the less affluent attend a non-religious school, but average street characteristics are assigned to the pupil postcodes, this would lead to a downward bias on $\%religsch_k$.

$$\begin{aligned} Y_{ijk.t} = \%religsch_k.KS4.\delta_1 + religsch_{jk}.KS4.\delta_4 + Z_{jk}.KS4.\delta_5 \\ + W_k.KS4.\delta_6 + X_{ijk}.KS4.\delta_7 + KS4.\delta_8 + \xi_{ijk} + \epsilon_{ijk.t} \end{aligned} \tag{8.2}$$

In the third specification, the competition effect is decomposed into an effect experienced by those attending religious schools and one for those attending secular schools. Again, sorting based on unobserved characteristics within areas mean there is a greater risk that the coefficients estimated are biased, relative to those estimated in equation 8.1. This specification is intended to reflect the possibility that the competition effect is not symmetrical in terms of its effect on pupils in the secular and religious sectors. In other words, the competition gain from attending a religious school in an area with many religious schools is not the same as the competition-related gain from attending a secular school in the same area.

$$\begin{aligned} Y_{ijk.t} = \%religsch_k.KS4.\delta_1 + \%religsch_k.religsch_{jk}.KS4\delta_3 \\ + religsch_{jk}.KS4.\delta_4 + Z_{jk}.KS4.\delta_5 \\ + W_k.KS4.\delta_6 + X_{ijk}.KS4.\delta_7 + KS4.\delta_8 + \xi_{ijk} + \epsilon_{ijk.t} \end{aligned} \tag{8.3}$$

In the fourth specification, the competition effect of religious schools on non-religious schools is once again decomposed into an effect on non-religious schools for whom a religious school is one of the three nearest secondary schools ($closereligsch_{jk}$), and those for whom there are no religious schools nearby in the county and so are likely to be under less competitive pressure. The purpose of this specification is to identify a more disaggregated competition effect, thus reducing the risk of type II errors posed by measuring competition levels over the large ancient county. However, there will be a sorting of pupils on unmeasured characteristics within counties, and if this sorting is in some way correlated with the geographical supply of religious schools, coefficients in this specification will be biased.

$$\begin{aligned} Y_{ijk.t} = {} & \%religsch_k.KS4.\delta_1 + \%religsch_k.closereligsch_{jk}.KS4\delta_2 \\ & + \%religsch_k.religsch_{jk}.KS4\delta_3 + religsch_{jk}.KS4.\delta_4 + Z_{jk}.KS4.\delta_5 \\ & + W_k.KS4.\delta_6 + X_{ijk}.KS4.\delta_7 + KS4.\delta_8 + \xi_{ijk} + \epsilon_{ijk.t} \end{aligned} \tag{8.4}$$

In order for the estimation of whole area effects to be valid, the presence of religious schools should not cause pupils to sort *across* areas. The models are estimated using ancient (1851) counties as the areal unit of analysis because this is the area for which the historical instrument (described later in this section) is available. There are 39 ancient counties in England, meaning they are much larger areas than the 150 modern local authorities. This increases the likelihood of the validity of the assumption that historic counties are contained markets without families re-locating across counties based on school types and quality. However, this is at the expense of possible aggregation bias on the measurement of exposure to religious schools.

Thirty-two equations are estimated in total: these four specifications are repeated across the four subject outcomes described in the data section for religious school competition, and this is repeated for Catholic school competition only. Competition from Catholic schools are of separate interest for two reasons. First, they represent two-thirds of all religious secondary schools, and are a more homogenous body of schools than the group of all religious secondary schools. So, it is possible that they exert a more consistent type of competitive effect on their neighbours. We have no reasons to believe that the effect of Catholic schools is likely to be similar to the effect of other faith schools. However, many Catholic schools will not be a realistic choice for the majority of parents because the religiosity requirements tend to be onerous and fewer families are able to claim some sort of relationship with the Roman Catholic Church, than they are with the established Church of England. The second reason is that it is only possible to find instrumental variables that provide an exogenous source of variation in Catholic schools, so estimation of the spillover effect of these schools is estimated including, and excluding, an IV for comparability.

8.3.1 Instrumenting Catholic school supply

There may be an endogeneity problem because the supply of places at religious schools is believed to be endogenous to current demand for religious schools, which is in turn related to the current religious population in the county and the quality of neighbouring Community schools. Both of these might independently enter the education production function for the county schooling system, thus biasing the estimates of the effect of religious schools. For example, if Catholic schools expand in response to poor quality Community schools, this would lead to a downward bias on $\%RCsch$. On the other hand, if Catholic schools expand to accommodate a larger Catholic population in the county and if high-religiosity families have characteristics that mean their children are academically successful (regardless of school attended), this would lead to an upward bias on $\%RCsch$.

This paper uses a source of variation in the supply of Catholic schools, $\%RCsch$, that can be argued to be exogenous to current demand for Catholic schooling. This variation comes from survey data on the historic sizes of the religious populations in England. It can be shown that the geographical distribution of historic Catholic populations predicts $\%RCsch$, conditional on the current RC population and other county control variables, i.e.:

$$\%RCsch_k = \alpha_0 + \alpha_1 \%RC1931_k + \alpha_2 \%RC2005_k + \alpha_3 Z_k + v_k \tag{8.5}$$

The F-value on the instrument for this first-stage is 20.81. The instrument is applied to the variable $\%RCsch_k$ as a first stage to the main pupil fixed effects specification, as set out in equation 8.1. The size of the Catholic population between 1851 and 1961 is sourced from survey data and measured once every ten years in each ancient county, with details of the survey described in Chapter 2. It makes relatively little difference to the results which year of data is used to predict $\%RCsch$. The size of the RC population in 1931 ($\%RC1931$) is chosen as an instrument because it is strongly correlated and it immediately precedes the opening of most RC secondary schools.

The claim of this identification approach is that the county-level controls are sufficient to meet the excludability restriction. These include the current proportion of Catholic church attendees in the area, and this is a crucial control variable given the absence of religious affiliation of the family in the pupil-level data (and therefore an inability to control for direct family religiosity effects). Unlike the data used by US researchers, there is not a good source of religious adherence by family that can be matched to NPD, and in any case it is not clear how valid a family's self-response to this question is likely to be. As has been pointed out by others, families who send their children to

religious schools may be more likely to claim religious adherence than those who do not, which would in itself be a source of bias (Neal, 1997). In addition to the current church attendance indicators for the county, a full set of socio-demographic indicators are also included to avoid confounding correlations between the type of areas (i.e. industrial and urban) that Catholics historically lived and the demographic characteristics of these areas that persist today.

8.4 Data

Data for this chapter are principally drawn from the National Pupil Database of school leavers at age 16 in 2005, further details of which can be found in Chapter 2. Four measures of pupil achievement are used in the chapter – test scores in English, Maths, Science and an aggregated test score. These are measured at Key Stage 2 (age 10/11) and Key Stage 4 (age 15/16), as shown in Table 8.5. The total subject scores (an average of attainment in the three subjects for KS2; the best 8 grades at GCSE for KS4) is transformed to a z-score to ease interpretation of comparisons across the Key Stages. Key Stage 3 test scores are also available, but are not used. They are difficult to interpret because they take place mid-way through secondary school and are therefore not high stakes exams from the perspective of the school or the pupils. A school that performs well on KS3 tests will boost their position in KS2-KS3 value-added score, but risks making it harder to achieve success in KS3-KS4 value-added tables.

Table 8.5: Key pupil test score variables

	Mean	Std. Dev.	Min.	Max.
KS2 English score	4.381	0.848	0.000	6.000
KS2 maths score	4.380	0.876	0.000	6.000
KS2 science score	4.665	0.698	0.083	6.000
KS2 total z-score	0.006	0.993	−5.712	2.055
GCSE English score	4.572	1.852	0.000	8.000
GCSE maths score	4.307	1.949	0.000	8.000
GCSE science score	4.240	1.999	0.000	8.000
GCSE capped z-score	0.012	0.991	−2.758	2.367

Complete cases summary statistics. Number of pupils=546,133

The standard set of pupil control variables are listed in Table 8.6. As in previous chapters, NPD is supplemented with data on the level of deprivation within the child's street, and 57 ACORN indicators of household type.

In a small number of specifications in this chapter, variables are included to indicate the type of school the child attends. These variables are listed in Table 8.7.

Table 8.6: Key pupil-level control variables

	Mean	Std. Dev.	Min.	Max.
Female	49.6%			
FSM	12.9%			
SEN statement	2.4%			
SEN action	9.2%			
SEN action plus	4.1%			
English not mother tongue	7.4%			
Ethnicity white British	83.3%			
Postcode deprivation (IDACI)	0.208	0.174	0.003	0.993
Postcode deprivation (IMD)	22.501	16.380	0.590	86.360

N=546,133. Controls also include 14 age, 14 ethnicity and 57 ACORN indicators

Table 8.7: Key school-level control variables

	Mean	Std. Dev.	Min.	Max.
Religious school	15.0%			
Catholic school	9.5%			
Girls school	6.2%			
Boys school	4.4%			
School % FSM	14.1%	12.4%	0.0%	87.0%
School % English not mother tongue	8.5%	16.3%	0.0%	100.0%

N=3,103, weighted for school size

The ancient county level variables list in Table 8.8 are sourced from several surveys. First, a large range of pupil-level characteristics are aggregated up to the ancient county level. Second, indicators of levels of church attendance by religious denomination are included from the 2005 English Church Census. Third, the proportions of the population identifying themselves as belonging to a religious group are included from the 2001 Census of Population.

Table 8.8: Key county-level control variables

	Mean	Std. Dev.	Min.	Max.
County % FSM	13.3%	5.0%	2.2%	2.5%
County % SEN statement	2.4%	0.6%	0.5%	4.6%
County % asian Indian	2.2%	2.9%	0.0%	13.6%
County % asian Pakistani	2.3%	2.0%	0.0%	6.3%
County % asian Bangladeshi	0.9%	1.4%	0.0%	6.1%
County % black African	1.5%	2.5%	0.0%	8.9%
County % black Caribbean	1.5%	2.1%	0.0%	7.3%
County % white British	81.5%	13.9%	36.0%	97.8%
County % English not mother tongue	8.7%	9.4%	0.4%	41.4%
County % Church attendance	6.2%	1.0%	3.8%	8.3%
County % CofE attendance	1.7%	0.4%	1.3%	2.8%
County % RC attendance	1.8%	0.7%	0.8%	3.0%
County % Christian (2001 census)	72.2%	6.1%	53.5%	83.6%
County % Hindu (2001 census)	1.1%	1.7%	0.0%	7.1%
County % Muslim (2001 census)	3.3%	3.0%	0.0%	12.0%
County % Jewish (2001 census)	0.4%	0.6%	0.0%	2.3%
County % Sikh (2001 census)	0.7%	0.8%	0.0%	3.0%
County % No religion (2001 census)	14.1%	2.6%	9.1%	18.5%
County average KS2 score	4.471	0.049	4.345	4.637
County average deprivation (IDACI)	0.210	0.054	0.065	0.332
County average deprivation (IMD)	22.687	6.651	7.211	33.020

N=39, weighted for county size

8.5 Results

8.5.1 All religious schools

This section analyses the results from the pupil fixed-effects test score growth model of the area-wide effect of religious schools. The total effect of all religious schools is analysed separately from the effect of Catholic schools. These first results are shown in Table 8.9. Results for all four subjects are discussed simultaneously.

The first specification shows that areas with more faith schools make greater educational progress in the all-subject score and in science (but not in English and maths).

The coefficient of 0.163 on the all-subject score indicates that the average child in an area with 25 per cent of pupils in religious schools would achieve 4% of a standard deviation improvement in overall GCSE performance compared to the same child in an area with no religious schools. The 0.418 coefficient on the science score indicates that the typical child in an area with 25 per cent of pupils in religious schools would do one-tenth of a grade better at GCSE science compared to living in an area with no religious schools. In other words, these effects are very small (though statistically significant) and there are no measurable effects for English and maths. The finding can be interpreted as unbiased estimates of these whole area effects provided there is not sorting across areas based on unmeasured characteristics that are correlated with $\%religsch$. The whole area effect could indicate some effect of competition from religious schools or could be entirely attributable to the differential effectiveness of religious schools, with pupils in non-religious schools in these areas doing no better at secondary school.

Specification (2) attempts to separate religious school effectiveness from a competition effect. Test score growth in every subject is superior in religious schools and the magnitude of estimates is consistent at just over one-tenth of a GCSE grade. However, inclusion of this variable produces inconsistent estimates of a competition effect, with it appearing positive and significant for science and the all-subject score, and negative and significant for English and maths. The problem here arises that if religious schools cause within-area sorting based on unmeasured characteristics, the coefficient on the competition effect is biased downwards.

The third specification analyses whether the competition effect appears to be symmetrical. It asks whether religious schools in areas where there are many religious schools do better than non-religious schools in these areas, and also whether they do better than stand-alone religious schools. The estimates show that it is not possible to reject the hypothesis that the non-religious and religious competition effects are equal, except for in maths and science where religious schools do worse than non-religious schools as a result of being in areas with many religious schools. Indeed, by adding δ_1 and δ_3, it can be seen that religious schools in high religious schooling areas suffer (except in the all-subject score) compared to stand-alone religious schools. In these circumstances they must compete for religious families, but this does not appear to induce them to work harder at improving exam results. Instead, they appear to suffer, possibly because the oversupply of religious school places makes the selection of a desirable intake less possible.

Specification (4) separates the effect of being a non-religious school located in a county with many religious schools in general from the effect of actually having one of these schools close by (nearest three schools). There is a positive and statistically significant

Table 8.9: Pupil fixed effects test score growth (religious schools)

All subjects z-score:	(1)	(2)	(3)	(4)
(δ_1) %religsch.KS4	0.163***	0.142**	0.146**	0.164**
	(0.048)	(0.049)	(0.049)	(0.051)
(δ_4) religsch.KS4	–	0.018***	0.023*	0.024**
		(0.004)	(0.009)	(0.009)
(δ_3) %religsch.religsch.KS4	–	–	−0.026	−0.049
			(0.047)	(0.051)
(δ_2) %religsch.closereligsch.KS4	–	–	–	−0.022
				(0.019)
R-sq (within)	0.091	0.091	0.091	0.091
Variance attributable to ξ_i	0.744	0.744	0.744	0.744
English score:	(1)	(2)	(3)	(4)
(δ_1) %religsch.KS4	−0.105	−0.288***	−0.271***	−0.311***
	(0.081)	(0.080)	(0.081)	(0.085)
(δ_4) religsch.KS4	–	0.128***	0.150***	0.147***
		(0.006)	(0.015)	(0.015)
(δ_3) %religsch.religsch.KS4	–	–	−0.126	−0.078
			(0.078)	(0.084)
(δ_2) %religsch.closereligsch.KS4	–	–	–	0.049
				(0.032)
R-sq (within)	0.212	0.226	0.226	0.226
Variance attributable to ξ_i	0.605	0.605	0.605	0.605
Maths score:	(1)	(2)	(3)	(4)
(δ_1) %religsch.KS4	0.002	−0.182*	−0.155	−0.313***
	(0.083)	(0.082)	(0.083)	(0.087)
(δ_4) religsch.KS4	–	0.120***	0.156***	0.144***
		(0.006)	(0.015)	(0.016)
(δ_3) %religsch.religsch.KS4	–	–	−0.207**	−0.013
			(0.080)	(0.086)
(δ_2) %religsch.closereligsch.KS4	–	–	–	0.194***
				(0.032)
R-sq (within)	0.212	0.228	0.228	0.228
Variance attributable to ξ_i	0.624	0.625	0.625	0.625
Science score:	(1)	(2)	(3)	(4)
(δ_1) %religsch.KS4	0.418***	0.204*	0.241**	0.121
	(0.093)	(0.093)	(0.093)	(0.098)
(δ_4) religsch.KS4	–	0.139***	0.189***	0.179***
		(0.007)	(0.017)	(0.018)
(δ_3) %religsch.religsch.KS4	–	–	−0.282**	−0.134
			(0.090)	(0.097)
(δ_2) %religsch.closereligsch.KS4	–	–	–	0.148***
				(0.037)
R-sq (within)	0.268	0.285	0.285	0.285
Variance attributable to ξ_i	0.524	0.525	0.525	0.525

N(pupils)=343,936; N(schools)=3,108; N(county)=39.

Note: ***=sig. at 0.1%; **=sig. at 1%; *=sig. at 5%.

effect for maths and science of being a non-religious school located close to a religious school compared to being a non-religious school elsewhere in the same county. If this is interpreted as a competition effect, it suggests that non-religious schools are only threatened by religious schools who are located close by. This seem plausible because, although religious schools do recruit from a wide geographical area, it is possible that the parents who make very long journeys to religious schools each day are the higher religiosity families, and so non-religious schools are not able to compete for these pupils.

There are clearly inconsistencies in findings across subject areas, which are difficult to explain. For example, the estimates on English are more consistently negative than for other subjects. The models are re-estimated on the sub-sample of only White British, English mother tongue pupils to examine whether the distribution of ethnic minority pupils across the country is responsible for this result, but the same pattern is found. It is true that achievement in English is more strongly influenced by families, rather than schools, and this may explain why the institutional environment does not appear able to positively affect achievement in this subject.

8.5.2 Catholic schools

Table 8.10 shows the same set of regressions for Catholic schools only. Because these constitute two-thirds of all religious secondary schools the results are very similar, but are presented with the addition of the instrumental variable specification. However, there is even less consistent evidence here for a competition effect from Catholic schools. The first specification shows that the presence of Catholic schools in an area is associated with significant slower progress in English and somewhat faster progress in science (with no significant effect on the all-subject and maths scores). These potential competition effects are little changed by the inclusion of the Catholic school dummy, although faster progress in maths now becomes marginally significant.

Specification (3) also shows that non-Catholic schools do better than Catholic schools as a result of being located in an area with many Catholic schools. Again, this is consistent with the thesis that Catholic schools suffer from oversupply and thus a deterioration in pupil quality. Finally, specification (4) shows that non-Catholic schools that are located close to Catholic schools do benefit more from being in an area with many Catholic schools than other non-Catholic schools in the county do.

The IV specification instruments $\%RCsch.KS4$ in specification (1) using $\%RC1931$. It produces a consistently negative and large effect of attending a secondary school in an area with many Catholic schools. The magnitude suggests that pupils in areas with 25 per cent of pupils in Catholic schools do half a grade worse in England and science and

Table 8.10: Pupil fixed effects test score growth (Catholic schools)

All subjects z-score:	(1)	(IV)	(2)	(3)	(4)
(δ_1) %RCsch.KS4	0.005	−1.063	−0.011	0.015	−0.003
	(0.106)	(36.169)	(0.106)	(0.107)	(0.109)
(δ_4) RCsch.KS4	–	–	0.016***	0.049***	0.048***
			(0.004)	(0.010)	(0.010)
(δ_3) %RCsch.RCsch.KS4	–	–	–	−0.256***	−0.236**
				(0.069)	(0.073)
(δ_2) %RCsch.closeRCsch.KS4	–	–	–	–	0.022
					(0.027)
R-sq (within)	0.091	0.090	0.091	0.091	0.091
Variance attributable to ξ_i	0.744	0.743	0.744	0.744	0.744
English score:	(1)	(IV)	(2)	(3)	(4)
(δ_1) %RCsch.KS4	−0.828***	−1.996***	−0.827***	−0.796***	−0.879***
	(0.177)	(0.235)	(0.176)	(0.176)	(0.180)
(δ_4) RCsch.KS4	–	–	0.124***	0.164***	0.160***
			(0.007)	(0.016)	(0.017)
(δ_3) %RCsch.RCsch.KS4	–	–	–	−0.310**	−0.215
				(0.114)	(0.121)
(δ_2) %RCsch.closeRCsch.KS4	–	–	–	–	0.104*
					(0.045)
R-sq (within)	0.213	0.213	0.226	0.226	0.226
Variance attributable to ξ_i	0.605	0.603	0.605	0.605	0.605
Maths score:	(1)	(IV)	(2)	(3)	(4)
(δ_1) %RCsch.KS4	0.355	−0.729***	0.396*	0.453*	0.171
	(0.182)	(0.242)	(0.180)	(0.181)	(0.184)
(δ_4) RCsch.KS4	–	–	0.093***	0.166***	0.152***
			(0.008)	(0.017)	(0.017)
(δ_3) %RCsch.RCsch.KS4	–	–	–	−0.569***	−0.249*
				(0.117)	(0.124)
(δ_2) %RCsch.closeRCsch.KS4	–	–	–	–	0.351***
					(0.046)
R-sq (within)	0.212	0.211	0.228	0.228	0.228
Variance attributable to ξ_i	0.624	0.623	0.625	0.625	0.625
Science score:	(1)	(IV)	(2)	(3)	(4)
(δ_1) %RCsch.KS4	0.467*	−2.226***	0.519*	0.575**	0.349
	(0.205)	(0.272)	(0.203)	(0.203)	(0.208)
(δ_4) RCsch.KS4	–	–	0.115***	0.187***	0.176***
			(0.009)	(0.019)	(0.019)
(δ_3) %RCsch.RCsch.KS4	–	–	–	−0.560***	−0.303*
				(0.132)	(0.140)
(δ_2) %RCsch.closeRCsch.KS4	–	–	–	–	0.280***
					(0.052)
R-sq (within)	0.268	0.267	0.285	0.285	0.285
Variance attributable to ξ_i	0.524	0.524	0.525	0.525	0.525

N(pupils)=343,936; N(schools)=3,108; N(county)=39.

Note: ***=sig. at 0.1%; **=sig. at 1%; *=sig. at 5%.

one fifth of a grade worse in maths. However, this is a Local Average Treatment Effect (LATE) estimate for the effect of Catholic school presence that resulted from large historical Catholic populations, relative to today. This might offer reasons as to why these estimates are negative. First, in these areas there is now a significant oversupply of Catholic schools, relative to the size of the Catholic population, and so the Catholic schooling sector is less likely to achieve an advantaged intake, and is therefore less desirable to parents. Non-Catholic schools may therefore not perceive the competitive threat from this sector to be particularly high. Second, areas with historically high levels of Catholics compared to current levels (e.g. London and the West Midlands) have continued to be areas where new immigrants have entered the country and settled, and so schools in these areas may face disadvantages and challenges that are captured in the LATE estimates.

8.6 Discussion

Religious secondary schools impact on all other schools in the area because they enable genuine choice of school for a group of parents, which produces a high degree of sorting in the local educational market. This chapter has examined whether this sorting impact spills over to the educational achievement of pupils who attend neighbouring schools, whether through school effort induced by competition for pupils, or changes in peer groups induced by sorting. Since religious secondary schools have been a long-term historical feature of the educational landscape in England, any competition effects identified might help inform policy-makers as to the long-term effects of enabling choice and competition, which may be very different to short-run responses by schools to market changes.

The chapter combined two estimation strategies previously used in the literature to attempt to identify these effects without the confounding influences of unmeasured pupil characteristics. The pupil fixed-effects test score growth models relied on the supply of religious schools in an area being uncorrelated with average unmeasured pupil characteristics that might independently enter the production function. The addition of the instrumental variable required the levels of the Catholic population in 1931 to be excludable from the education production function, once the modern-day religious and social characteristics of the county were accounted for.

Overall, the regressions fail to find a consistently positive (or negative) effect of religious schools on overall area-wide educational performance. However, the pupil fixed-effect regressions provide some evidence of positive area-wide effects resulting from the presence of religious schools and suggest that there are benefits to attending a non-religious school located close to a religious school, rather than one located further away in the

same county. However, with the within-area comparisons there is a high risk that the estimates are confounded by pupil sorting. The instrumental variables approach suggests that areas with large numbers of Catholic schools due to a large historical Catholic population do significantly worse in terms of GCSE performance. The most likely explanation for this negative estimate is that these areas have continued to house new immigrants to the country and that this presents continuing challenges to schools. Alternatively, the negative estimate might be due to the nature of the oversupply of Catholic school places in many of these areas (Grace, 2002, reports that large numbers of Catholic schools in areas such as Birmingham have many non-Catholic pupils). The IV (given the modern day Catholic population controls) estimates the effect of significant oversupply of Catholic schools, which means the Catholic sector struggles to fill its capacity, has a lower social and ability profile of intake, and thus presents a low competitive threat to non-religious schools.

There are three possible explanations for the lack of a consistent positive competition effect as the result of religious schools being present in an area. The first possibility is that religious and non-religious schools do not actually compete for pupils. In other words, high levels of pupil mobility and sorting in an area is not sufficient to suggest that competition between schools for pupils is actually taking place. This would be true if they essentially operate in separate markets, with each sector recruiting from different primary schools. There is some evidence for this in the case of Catholic schools, which use feeder school admissions criteria to segment the market. This might explain why the results for Catholic schools are less positive.

The second explanation is that headteachers in non-religious schools do feel genuinely threatened by the presence of local religious schools, but they do not (or cannot) respond to this threat. The most likely reason for this is that they do not possess the means to significantly influence effort exerted on the part of their classroom teachers so the threat does not translate into improve GCSE performance (i.e. there is a principal-agent problem). Alternatively, they may find that, although the number and quality of pupils at their schools is being affected by a religious school's presence, given little spare capacity in the system their school's position is sustainable so no effort response is necessary for survival.

The final explanation is that competition is actually muted, rather than increased, by the presence of faith schools, because they allow the system to become stratified. This stratification then provides schools with a disincentive to focus effort on improving test scores because marginal changes in effort cannot affect a school's league table position. There is a positive correlation between the level of FSM and top ability segregation and the number of religious or Catholic schools in ancient counties (e.g. $\rho = 0.52$ for $\%RCsch$ and FSM segregation). Because of this stratification, religious schools are

also associated with an area-wide stretching of the ability distribution between KS2 and KS4 (there is a positive correlation of $\rho = 0.45$ between the increase in the county standard deviation in test scores from KS2 to KS4 and the proportion of pupils in religious schools).

This observation that many schooling systems have a tendency to become stratified in the long-run is one of the problems with operating a quasi-market. Thus, it is possible that relatively recent competition reforms in countries such as the US and Sweden might have identifiable short-run competition effects, but if the system is also stratifying, efficiency is likely to fall in the long-run.

Chapter 9

Conclusion

This thesis has presented new evidence on the effects of the English system of secondary school admissions on social welfare. In doing so it contributes to the quantitative literature on choice-based school admissions, pupil sorting and school competition. This conclusion will draw together the research findings from preceding chapters and discuss the implications of these results for policy-makers involved in the development of market-based reforms to schooling.

9.1 Parental choice of school

Choice of secondary school appears to be genuine and feasible for some households in many parts of England, although restricted school capacity imposes very real constraints on choice for almost all parents. Through the exploitation of recently available postcodes in the National Pupil Database, Chapter 4 showed that half of all English pupils do not appear to attend their nearest secondary school. However, most of this apparent choice relates to pupils who are almost certainly at a *de facto* neighbourhood school, or are at grammar or Voluntary Aided religious schools that were in place prior to the 1988 Education Reform Act. It places an upper bound on the number of pupils involved in sorting between non-faith comprehensive schools, whether voluntarily or involuntarily, at one-in-five.

The statistics quoted from Chapter 4 are not good estimates of potential or even exercised choice because attending a particular non-local school may not reflect the actualisation of a real preference. It is possible that the nearest school had oversubscription criteria that excluded the child, or that another non-local school was their true preference. Equally, a great many children attend their local school as their most preferred option, and so may be active choosers as much as those that travel further to school.

In theory, offering parents a choice of secondary school for their child should in itself be welfare enhancing. However, the problem with the implementation of choice in the English system is that it can be illusory because parents are invited to express a preference for a school, even where there is close to a zero probability that it would be prepared to accept their child. The *ex-post* realisation of this causes some parents to find the process dissatisfactory. Furthermore, even where genuine choice of school exists, having a choice between a poor school and a good school, where previously the only option was the good school does not increase well-being (Kelly, 2007). In the English system, with little creation of new capacity and parental preferences for school type that are reasonably homogenous, the satisfaction of one person's choice necessarily denies that of another, meaning that choice on its own may not be welfare improving.

The empirical literature on school choice in England has not yet succeeded in measuring how much *achievable choice* parents have, in terms of the size and nature of the choice set of schools who would actually accept their child, holding constant house location and the preferences of others. A first step towards calculating this type of measure would be to gain access to the stated preference forms submitted by parents of year six children and held by local authorities and other admissions authorities. This information could be combined with school oversubscription criteria to identify achievable choice sets for individual applicants. It is likely to reveal very large disparities in the levels of achievable choice, even between families who live in the same local authority, and there will almost certainly be a social class dimension to these inequalities because some oversubscription criteria discriminate between applicants based on the attributes of the family or the child.

This thesis has been able to show that schools are more segregated than the neighbourhoods in which they are located, confirming that where pupils are sorting themselves into a non-proximity school, it does tend to increase social and ability segregation between schools, relative to underlying residential segregation. This should be taken as an indication that processes allocating pupils to schools may in some way be inequitable with the result that high quality schooling is unevenly distributed across the social classes. This may be because low income families are financially constrained in their ability to make choices, or they are unable to meet the criteria to gain places at popular schools, or alternatively they may not be choosing to engage in the choice process. Regardless of which of these reasons dominates, we can reject the claim that choice policies have disproportionately benefitted low income children because their families were previously unable to afford homes close to popular schools, thus lowering school segregation. This lowering of school segregation through choice has not materialised because England did not start from a position of complete residential stratification, and in addition policies have given some schools both the means and the motivation

to recruit pupils with fewer problems and above average ability.

Although the thesis makes few conclusive statements about causes of school stratification, it is notable that religious, and to a lesser extent, Foundation schools have intakes that are more advantaged than their local neighbourhoods. These claims regarding the location of stratification in the secondary schooling system are consistent with a longitudinal analysis of FSM segregation that shows school segregation has not increased since the 1988 Education Reform Act (e.g. Gorard et al., 2003). Chapter 4 notes that most post-residential sorting is attributable to the admissions policies of grammar and religious schools that were in place prior to the 1988 Education Reform Act. Indeed, given that 1990 to 2005 was an era of rising pupil numbers and falling spare capacity, it is possible that achievable choice did not actually rise over this period. Thus, secondary schools may be no more segregated than they were in the 1980s, but parents may also have little additional choice.

The findings of this thesis should certainly not be taken as evidence that segregation is the price that must be paid for enabling free parental choice and facilitating competition between schools. It is true that there exists a social class gradient in the capacity (and desire) of parents to engage in the school choice process, but this has been exacerbated by the complexities of the English choice system which continues to sanction variation in admissions procedures across state-funded schools and allows oversubscription criteria so complex that it is impossible for a family to assess the probability of achieving a place at their desired school. Policies to simplify admissions procedures may in themselves be more equitable, and in addition simplification may encourage low income families to engage with the system. However, this claim is somewhat speculative because our limited success in modelling parental choice means we lack the capability to predict whether policies directed at admissions simplification, the employment of choice advisors or transport subsidies, for example, would best serve to improve the educational opportunities of children from poor families. Given existing oversubscription criteria used by schools, it remains possible that engagement with choice is unproductive if the overwhelming constraint on the poor is their severely reduced achievable choice set, because their children are less likely to pass academic selection tests or successfully demonstrate religious adherence.

9.2 School admissions reforms

The main policy recommendations that arise from this thesis relate to changes in sanctioned secondary school admissions rules to improve fairness in the system. Fairness would seem to be a crucial goal because school admissions reforms (given capacity constraints), do not usually facilitate choice, but instead have the effect of transferring a

particular set of educational opportunities from one child to another. Chapter 5 shows that own-admissions schools have intakes that are more advantaged than Community schools, even when the characteristics of local neighbourhoods are taken into account. If the goal of admissions reforms is to lower the level of social stratification between schools, this thesis concludes that reforms to the admissions policies of religious, Foundation and grammar schools should be an important starting point.

Religious schools across denominations and regions have intakes that are, on average, significantly more advantaged than the neighbourhood in which they are located. The data presented in Chapter 5 suggest that if we take a Community school and a religious school located in neighbourhoods with the same demographics, the religious school might have as many as 50 per cent more top ability pupils than the Community school does. Given that the religious background and practices of the family are a social phenomenon, it is not surprising that religious adherence criteria are correlated with social class. However, there remains an unanswered question as to why the social composition of religious schools appears to be so much more advantaged than that of the church-going population.

The current measurement of religious adherence on a 'continuum' justifies the collection of family background data, giving religious schools the means to socially select pupils, should they wish to do so. It would seem inherent that problems will always arise in this process, because some questions deemed as relevant to establish religious adherence – such as marital status and place of child's baptism – by their nature reveal information about the social background of the family. Furthermore, even without explicit cream-skimming taking place by religious schools, the complexity of their current admissions criteria may discourage low income families from applying, or alternatively they may apply but be less skilled at meeting a specific school's criteria for religiosity. Thus, it is possible that religious schools are cream-skimming inadvertently in the process of selecting by religious adherence. One way to simplify the admissions process for all families would be for the churches themselves to establish a nationally agreed binary criteria of 'religious adherence' that families are deemed to have either met, or not met. Once this is established, religious schools could then rely solely on the presence of a signature on a form from a religious leader to decide who has priority in the admissions process, so avoiding the need for the schools themselves to collect family background information.

Although Voluntary-Aided religious and Foundation schools could be argued to have similar means and motivations to cream-skim, Chapter 5 shows that the Foundation schools sector has an intake only marginally more advantaged than the neighbourhoods in which they are located. This assertion that most Foundation schools do not appear to be engaged in cream-skimming activities challenges the conclusions drawn by those

who simply measured the social and ability characteristics of these schools without accounting for their greater propensity to be located in more affluent parts of the country. This finding raises more questions than it answers about why schools who are given very clear incentives to cream-skim as the lowest cost route to raising their league table position, and therefore popularity, usually choose not to do so. The tentative suggestion made in Chapter 5 is that school leaders at these former LA-controlled schools may have a strong commitment to their local community and a notion of 'fair play', so will only introduce overtly selective criteria where they see it as essential for their own survival or where they perceive that other schools are cream-skimming, for example because there are grammar schools or other partially selective schools in the area.

However, a notable minority of these schools do use ability or aptitude tests to select up to 30 per cent of their intake, and this significantly impacts on the social composition of neighbouring schools. Removing this right to select, along with the selection that continues in the 164 grammar schools would substantially lower social stratification, but raises potential conflicts between the desire for greater educational opportunities for low income children and efficiency considerations regarding the education of the most able in society that are beyond the scope of this thesis. That said, one straightforward policy to reduce social stratification would be to remove the right of automatic entry to partially selective schools for the younger siblings of pupils who secure selective places, since they displace others who live closer to the school yet have themselves displayed no aptitude for the school's specialist subject. There is no clear rationale for allowing this policy to continue (it would be unthinkable for the younger siblings of grammar school pupils to be given automatic right of entry) and the policy has enabled a minority of 'comprehensive' schools to exclude almost all neighbourhood pupils.

There have been major changes to the school admissions code since the pupils under analysis in this thesis started secondary schools and there has not yet been an analysis of the impact of these changes on stratification. Chapter 5 demonstrates that explicit admissions policies are an important source of school stratification, which suggests that the 2007 School Admissions Code, combined with better adherence to the Code by schools, should help to balance intakes across the schooling system. Following the release of NPD for pupils starting secondary school in 2008, three surveys of school admissions policies (West et al. (2004) for 2001 entry; Coldron et al. (2008) for 2006 entry; and a forthcoming West survey for 2008 entry) can be used to measure how specific changes in a school's admissions policy, brought about by changes in legislation, alters that school's intake. This work is particularly important because it measures the direct relationship the policy lever of a national school admissions code, the implementation of the legislation by schools and the resulting composition of school intakes.

Despite significant legislative reform, it remains true that the objectives of an effective admissions system have remained largely implicit and there has not been enough analysis of what they might be (Coldron et al., 2008). However, one goal should be, and has been according to Government, that they the achieve fairness in school allocations. The concept of 'fair access' was enshrined within the 2007 Code's statutory legislation with Alan Johnson, then Secretary of State for Education, writing that the advent of the new Code would ensure that admissions procedures '*operate in a fair way that promotes social equity and community cohesion*' (Department for Education and Skills, 2007b, page 7). However, the word 'fairness' has arguably been narrowly defined by Government and the Schools Adjudicator. The Adjudicator, in a ruling on the new Brighton and Hove admissions system, judged the requirement of the Code that policies must not disadvantage, directly or indirectly, a child from a particular social or racial group, in the very narrow sense that it is only prohibited to design policies that excludes certain groups if the policy explicitly sets out to do so (Eastwood and Turvey, 2008). This means that new admissions policies can cause an increase in social stratification in an area, provided they do not explicitly intend to do so.

9.3 The proximity oversubscription criterion

Reforms to schools admissions criteria regarding religiosity and academic selection do have the potential to produce very real reductions in social stratification between schools and therefore modestly improve the average peer group experienced by children from poor backgrounds. The potential reductions in school stratification by these types of reforms are potentially very large in the very areas, such as Haringey and Wandsworth, where school segregation is seen to be a problem. However, Chapter 4 shows these policies are marginal in many parts of the country compared to the huge impact that the use of catchment areas and proximity oversubscription criteria have on pupil sorting. Part of this residential stratification is directly caused by the use of these criteria, although a large proportion would persist regardless of the schooling system. We do not currently have good estimates for the amount of residential stratification that is attributable to parents deliberately relocating their family to access a particular catchment area, although a replication of methods in Chapter 4 using postcodes from age 5 through to age 11 could indicate the amount of sorting that takes place during these years.

Systems involving lotteries and banding are potential reforms that would be consistent with achieving more socially integrated schooling. The allocation of places at all schools based on open lotteries for places among applicants would appear to be truly blind to the social background of the applicants. However, schools would still be stratified given that the housing market stratifies for reasons that are unrelated to school catchment

zones and the (time and money) costs of daily transport to and from schools would restrict choice for all. Also, by not giving families a guaranteed place at a local school, it would be a stressful process resulting in enforced long, or even infeasible, journeys for some families, with associated high costs that may have to be borne by the State.

A system of local authority-wide banding with a proximity oversubscription criteria within, say, five ability bands (i.e. similar to the old Inner London Education Authority (ILEA) scheme) would be more likely to achieve integrated schooling while minimising journey times. It is not inconsistent with a choice system (as many critics claim) since places within each band can be allocated based on parental preferences for schools. However, this type of scheme has several potential problems. First, the ILEA experience suggests that it can lead to gaming by parents who encourage their children to score badly on banding tests if they live in affluent areas and suspect it will increase the probability of a place at their favoured school. Second, if schools with intakes that are representative of children in the LA are unacceptable to certain parents they will seek an alternative school in the private sector or in another LA. Given this activity will be more prevalent among high income households, this in turn further lowers the ability distribution of the home LA, making it unacceptable to a further set of parents. Thus, the system can be unstable and lead to *across* LA stratification. The extent to which this happens depends on the size of LAs and how easy these cross-LA journeys are. Third, achieving balanced intakes is very difficult in large or rural LAs because it might involve transporting pupils long distances. Finally, it potentially lowers incentives for schools to compete for pupils, compared to the current system, because schools can only compete for pupil numbers and not for quality. On the other hand, however, if intakes are truly balanced, school quality will be clearly revealed through simple league tables, which should direct parents to apply for the highest quality schools. Competition for pupil numbers can then be made more effective by mechanisms to increase capacity at popular schools.

Although banding has significant potential problems, it is consistent with some commitment to neighbourhood schooling, which has benefits that should not be understated. It gives certainty of future transitions for parents, does not require additional testing at age 10/11, allows pupils to progress to secondary school with their primary school friends and minimises the journeys pupils must make to school each day. Therefore the best feasible policy reforms are likely to retain aspects of the neighbourhood schooling system with policies to enhance choice and equity. Areas of the US that have experimented with choice reforms have usually retained neighbourhood priority systems (e.g. a 3 mile walk zone around schools in Boston and defined catchment areas in Charlotte), but give parents the opportunity to express a preference for alternative schools. Priority for all out-of-catchment pupils is allocated by lottery in most systems. A more radical version of this approach would assign out-of-catchment places based

on child's ability, with priority given to children who balance the intake of the school to make it more representative of the area.

In the UK, as the pupil population shrinks, there exists a narrow window of opportunity to reform school admissions without disenfranchising families who have bought homes in the expectation of achieving a school place. Usually falling rolls produce an expansion in the geographical area a school recruits from. However, rather than allowing new streets to become part of the *de facto* school catchment, the spare places produced by falling rolls could be set aside and allocated based on a lottery or some compensating intake principle.

9.4 Predicting the effects of future policy reforms

Although researchers are able to make generalised claims about potential effects of these admissions policy reforms, they are unable to predict the magnitude of effects on school sorting, residential sorting, house prices, and so on. This is because the empirical literature, to which this thesis contributes, confines itself to documenting the types of school admissions procedures that are currently being used by schools, associations between admissions policies and stratification, differences in the propensities of pupils to attend particular schools, and areas of the country where pupil sorting appears to be a problem (e.g. Burgess et al., 2007; Chamberlain et al., 2006; Coldron et al., 2008; West et al., 2004).

Chapter 3 documented the huge gap that exists between the theoretical models that might be most appropriate for building predictive models of sorting and the current state of the literature. It proposes that an agent-based modelling framework might be more appropriate than a general equilibrium solution given that households display myopic behaviour and have financial and psychic attachments to particular areas, so as a result housing markets can take generations to adjust to policy changes. The sociological choice literature is capable, in part, of informing decisions about how to specify household utility functions and behaviours. However, it is not a quantitative empirical literature and so can play no part in parameterising the model.

There have been two approaches to parameterising school choice models by economists. The first is to find an extremely rich dataset, with full micro information on neighbourhood and school qualities, household characteristics and house prices, as Bayer and McMillan (2005) did in San Francisco. A structural equation model can then be used to estimate values for the parameters in the model. However, this approach can only be used in an equilibrium modelling framework and is invalid if the theoretical model is incorrectly specified. The alternative (and more viable) approach is to seek estimates for the values of particular parameters in isolation, often by finding

sources of exogenous variation in the data, as the school district boundary estimates of willingness-to-pay for school quality claim to do (e.g. Black, 1999).

9.5 School competition policies

Competition between schools for pupils provides a potential route by which standards can be raised across the secondary education system, regardless of how choice affects sorting. Both Foundation and religious secondary schools appear to enable genuine choice of school for a group of parents, which results in them recruiting from a much wider catchment area than their direct neighbourhood and produces a high degree of sorting in the local educational market. It is likely that secondary schools adjacent to Foundation and religious schools do perceive them to be a competitive threat, but this apparent competition for pupils does not appear to translate into effort on the part of neighbouring schools to increase exam performance. The only potentially positive competition effect findings in this thesis are from the pupil fixed-effect regressions in Chapter 8 that suggest areas with many religious schools perform better in some subjects at GCSE, and that there are benefits to attending a non-religious school located close to a religious school, rather than one located further away in the same county. However, this second claim is based on within-area comparisons, so there is a high risk that the estimates are confounded by pupil sorting. No spillover effects of Foundation schools onto area-wide achievement are found in Chapter 7. Thus, overall this thesis fails to lend support to the assumption underlying market reforms that structures to encourage schools to compete for pupils are a route to improving standards.

The lack of identifiable competition effects is broadly consistent with existing quantitative literature on school competition in England, yet contrasts with positive competition effects that have been identified in other countries. This reflects the reality that the effect of policies is highly contingent on the circumstances in which they operate. It is also true that studies from the US and Sweden usually measure the short-run effect of large changes to the schooling market through the introduction of new Charter or state-funded private schools, compared to long-run effects analysed in this thesis. It is possible that these large scale disruptions to pupil allocations do force schools to compete for pupils, but as the new hierarchy of a local market is re-established on new terms, it is possible that the incentive or imperative to compete for survival once again subsides.

Variation in levels of school competition is hard to measure, and this could explain the null results found in this thesis and elsewhere. However, there are more likely to be real substantive reasons related to the institutional environment that explain

why encouraging competition through policies of school autonomy is not an effective route to raising standards. Regarding the null result found for religious schools, it is likely that many religious and non-religious schools do not actually compete for pupils. In other words, high levels of pupil mobility and sorting in an area is not sufficient to suggest that competition between schools for pupils is actually taking place. This is likely true for Catholic schools who may essentially operate in separate markets, using feeder school admissions criteria to segment the system. However, this type of non-competition explanation would not seem to be applicable to the former Grant-Maintained schools.

The second explanation is that headteachers in Community schools do feel genuinely threatened by the presence of local religious or Foundation schools, but they do not (or cannot) respond to this threat. The most likely reason for this is that they do not possess the means to significantly influence effort directed at test scores exerted on the part of their classroom teachers. Alternatively, they may find that, although the number and quality of pupils at their school is being affected by an autonomous school's presence, given little spare capacity in the system their school's position is sustainable, so no effort response is necessary for survival.

The final explanation is that competition is actually muted, rather than increased, by the presence of religious (and to a lesser extent Foundation) schools, because they allow the system to become stratified. This stratification then provides schools with little incentive to focus effort on improving test scores because marginal changes in effort cannot affect their league table position.

This thesis, as with most large-scale quantitative fixed-design research, fails to give an account for why a competition effect cannot be found. It would now seem that this literature is incapable of progressing through further quasi-experimental econometric studies. What is lacking is a detailed institutional understanding of schools and of the motivations and behaviours of headteachers, governors and classroom teachers. So, this body of quantitative research on school competition should be supplemented with surveys to try to identify why competition does not appear to be effective. The survey of headteachers by Levačić (2004) is one direction of research that is undoubtably helpful, but the collection of data on perceptions and practices of classroom teachers as well as headteachers would be more beneficial, for it is their behaviour that ultimately determines pupil achievement.

There are three types of policy interventions that might be successful at encouraging schools to compete through greater effort focussed on pupil achievement. The first would be policies aimed at incentivising classroom teachers to maximise pupil test scores. A natural route to implementing this would be through a teacher performance related pay scheme. However, these are political contentious and difficult to imple-

ment because they require the measurement of teacher and school quality, without the confounding influence of pupil background characteristics. Adnett and Davies (2005) suggest a more effective policy would be to implement schemes to encourage teachers or departments to compete for pupils within schools, possibly with additional pay and resources directed at departments that are successful in attracting pupils to study their subject. This type of scheme may be more effective in increasing effort, given there exists more variation in quality within schools than between schools. However, it suffers from similar problems to existing quasi-market policies, namely effort might be re-directed towards borderline pupils and departments may engage in cream-skimming of easier-to-teach pupils.

Secondly, admission policy reforms directed at lowering stratification might also be successful in encouraging competition through school effort focussed on achievement. Schools that are more closely matched in terms of pupil composition have more incentive to increase marginal effort directed at test scores in order to attract parents to the school. One of the key problems with operating a quasi-market appears to be a tendency for schools to become stratified in the long-run. Indeed this might explain why short-run competition effects can be found following the implementation of new schemes in the US and Sweden. However, if these reform policies also allow the system to stratify, even only slightly, efficiency is likely to fall back in the long-run.

Finally, an essential condition for competition to be an effective route to school improvement is the constant injection of new capacity at popular schools or new schools. The current system in England does have spare capacity, but it always quickly settles in the weakest schools in the local hierarchy, who present no competitive threat to other schools. Capacity at the best schools does not expand because it is not in their interests to allow the least popular schools to close. The introduction of new capacity can be introduced via several means including forcing or incentivising existing popular schools to expand, or opening new schools (whether maintained, independent not-for-profit or for-profit schools funded by quasi-vouchers). However, regardless of how this is done, the constant need to create new capacity means that true choice and competition have a large financial cost, so that unless the reforms succeed in significantly raising standards, they may lower the efficiency of the schooling system.

Without a deeper understanding of why competition is currently ineffective at raising standards, there is no way to judge which reforms are likely to improve school quality at lowest cost. However, given the substantial costs associated with implementing competition policies, they should be judged against policies that do not seek to encourage schools to compete. One problem with encouraging competition, particularly if it leads to cream-skimming activities, is that it rules out the possibility of collaboration between schools in the local market, which could lead to reductions in the speed of

dissemination of good practice, less (potentially risky) innovation by schools at the top of the hierarchy, and an inability to innovate by schools who are severely financially constrained as a result of falling rolls (Adnett and Davies, 2003). Thus a policy choice exists between promoting co-operation and promoting competition, which headteachers see as mutually exclusive conditions (Ribchester and Edwards, 1998). Adnett and Davies conclude that:

> Competition is more likely to promote short-run efficiency and co-operation is more likely to promote long-run dissemination. Whether competition or co-operation is more likely to promote effective innovation depends on the strength of market hierarchies, first-mover advantages, and the resources required for successful innovation.
>
> (Adnett and Davies, 2003, page 194)

However, the likely success of this alternative non-market strategy is entirely contingent on our beliefs about whether teachers, as public servants, can be persuaded to act as knights, innovating and collaborating in the interests of educational progress despite an absence of threat to school resources and survival (Le Grand, 2003). Arguably this perspective can be questioned, given that the period leading up to market-based reforms in schooling was not associated with great educational innovations and rising standards. Given the motivations and qualities of teachers today, there may be no alternative to making a system of school accountability, through published school standards and the opportunities and threats of the quasi-market, work more effectively.

9.6 Concluding remarks

The introduction of a quasi-market for school places has been central to government aims to improve school standards in England for twenty years. However, the findings of this thesis are rather equivocal on the social welfare benefits of quasi-market reforms. Choice does appear to be possible for many parents, and this can have value in itself. However, ability to access high-performing schools continues to be unevenly distributed by family background. This, combined with the continuing right of some schools, particularly grammar and religious schools, to select pupils based on ability and religious adherence has led to a system that appears to be more stratified than a neighbourhood schooling system is likely to be.

The stratifying tendency of current secondary school admissions in England results in a system that is inequitable, without measurable efficiency gains produced through effort induced by competition between schools for pupils. However, this should not be

taken as evidence that quasi-market reforms cannot benefit society; it simply reflects the specific way that the English schooling system is currently managed. Devising a more efficient and equitable system of school admissions for England is extremely complex. However, there are clear policies to reform school admissions that should be successful in lowering school stratification, thereby simultaneously increasing the extent to which school effort is likely to be directly rewarded with greater demand for places and also raising educational opportunities for children from socially disadvantaged families.

Bibliography

Abernathy, S. F. and Mackie, C. J. (2002). Skimming, social welfare and competition among public and private schools: an agent-based perspective. Paper presented at the annual meeting of the American Political Science Association, Boston, Massachusetts.

Adnett, N. and Davies, D. (2002). Education as a positional good: Implications for market-based reforms of state schooling. *British Journal of Educational Studies*, 50:189–205.

Adnett, N. and Davies, D. (2003). Schooling reforms in England: From quasi-markets to co-opetition? *Journal of Education Policy*, 18:393–406.

Adnett, N. and Davies, D. (2005). Competition between or within schools? Reassessing school choice. *Education Economics*, 13:109–121.

Adnett, N. and Davies, P. (1999). Schooling quasi-markets: Reconciling economic and sociological analyses. *British Journal of Educational Studies*, 47:221–234.

Ahlin, A. (2003). Does school competition matter? Effects of a large-scale school choice reform on student performance. *Deparment of Economics, Uppsala University Working Paper*, 2.

Allen, R., Burgess, S., and Windmeijer, F. (2008). Correcting segregation indices for small unit bias. Mimeo. University of Bristol.

Allen, R. and Vignoles, A. (2007). What should an index of school segregation measure? *Oxford Review of Education*, 33(5):643–668.

Allen, R. and West, A. (2009). Religious schools in London: School admissions, religious composition and selectivity. *Oxford Review Of Education*, forthcoming.

Altonji, J. G., Elder, T., and Taber, C. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *Journal of Political Economy*, 113:151–84.

Ammermueller, A. and Pischke, J.-S. (2006). Peer effects in European primary schools: Evidence from PIRLS. *NBER Working Paper*, 12180.

Anderson, L. (2000). Farewell to Grant Maintained status: The future of self-governing schools. *School Leadership and Management*, 20(3):371–385.

Angrist, J., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Society*, 91:448–458.

Armor, D. J. and Peiser, B. M. (1998). Interdistrict choice in Massachusetts. In Peterson, P. E. and Hassel, B. C., editors, *Learning from school choice*. Brookings Institution Press, Washington, DC.

Atkinson, A., Gregg, P., and McConnell, B. (2006). The result of 11 plus selection: An investigation into opportunities and outcomes for pupils in selective LEAs. *CMPO Working Paper*, 06/006.

Audit Commission (1996). *Trading places: The supply and allocation of school places*. The Audit Commission, London.

Ball, S. J. (1993). Education markets, choice and social class: The market as a class strategy in the UK and the USA. *British Journal of Sociology of Education*, 14:3–19.

Ball, S. J. (2003a). *Class strategies and the education market*. RoutledgeFalmer, London.

Ball, S. J. (2003b). The risks of social reproduction: The middle class and education markets. *London Review of Education*, 1:163–175.

Bankston, C. and Zhou., M. (2002). Social capital and immigrant children's achievement. *Schooling and Social Capital in Diverse Cultures*, 13:13–139.

Barber, M. (1994). *The making of the 1944 Education Act*. Cassell Education, London.

Bast, J. L. and Walberg, H. J. (2004). Can parents choose the best schools for their children? *Economics of Education Review*, 23:431–440.

Bayer, P. (2000). Tiebout sorting and discrete choices: A new explanation for socioeconomic differences in the consumption of school quality. Mimeo. Yale University.

Bayer, P., Ferreira, F., and McMillan, R. (2004). Tiebout sorting, social multipliers and the demand for school quality. *NBER Working Paper*, 11087.

Bayer, P. and McMillan, R. (2005). Choice and competition in local education markets. *NBER working paper*, 11802.

Belfield, C. R. and Levin, H. (2003). The effects of competition on educational outcomes: A review of US evidence. *Review of Educational Research*, 72(2):279–341.

Benn, C. and Chitty, C. (1996). *Thirty years on: Is comprehensive education alive and well or struggling to survive?* Penguin Books, London.

Bettinger, E. P. (2005). The effect of charter schools on charter students and public schools. *Economics of Education Review*, 24:135–147.

Bifulco, R. and Ladd, H. F. (2004). The impacts of charter schools on student achievement: Evidence from North Carolina. Sanford Institute Working Paper Series.

Bjorklund, A., Edin, P.-A., Fredriksson, P., and Krueger, A. (2004). Education, equality and efficiency - an analysis of Swedish school reforms during the 1990s. *IFAU report*, 1.

Black, S. E. (1998). Measuring the value of better schools. *Economic Policy Review*, 4:87–94.

Black, S. E. (1999). Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics*, 114:577–99.

Board of Education (1941). Education after the war. Green Book.

Bogart, W. T. and Cromwell, B. A. (1997). How much more is a good school district worth? *National Tax Journal*, 50:215–232.

Bogart, W. T. and Cromwell, B. A. (2000). How much is a neighbourhood school worth? *Journal of Urban Economics*, 47:280–305.

Böhlmark, A. and Lindahl, M. (2007). The impact of school choice on pupil achievement, segregation and costs: Swedish evidence. *IZA Discussion Paper*, 2786.

Bohte, J. (2004). Examining the impact of charter schools on performance in traditional public schools. *Policy Studies Journal*, 32(4):501–520.

Booker, K., Gilpatric, S., Gronberg, T., and Jansen, D. (2005). The effect of charter schools on traditional public school students in Texas: Are children who stay behind left behind? *National Center for Study of Privatization in Education Occasional Paper*, 104.

Bourdieu, P. (1986). *Distinction: A Social Critique of the Judgement of Taste*. Routledge, London.

Box, G. E. P. (1979). Robustness in the strategy of scientific model building. In Launer, R. L. and Wilkinson, G. N., editors, *Robustness in statistics*. Academic Press, New York.

Bradley, S., Crouchley, R., Millington, J., and Taylor, J. (2000). Testing for quasi-market forces in secondary education. *Oxford Bulletin of Economics and Statistics*, 62(3):357–390.

Bradley, S., Draca, M., and Green, C. (2004). School performance in Australia: Is there a role for quasi-markets? *Australian Economic Review*, 37:271–86.

Bradley, S., Johnes, G., and Millington, J. (2001). The effect of competition on the efficiency of secondary schools in England. *European Journal of Operational Research*, 135:545–568.

Bradley, S. and Taylor, J. (2007). Diversity, choice and the quasi-market: An empirical analysis of secondary education policy in England. *Lancaster University Management School Working Paper*, 038.

Breen, R. and Goldthorpe, J. (1996). Explaining educational differentials: Towards a formal rational action theory. *Rationality and Society*, 9:275–305.

Brierley, P. (2006). *Religious trends 6: Pulling out of the nosedive*. Christian Research, Swindon.

Brown, P., Halsey, A., and Lauder, H. (1997). The transformation of education and society: An introduction. In Halsey, A., Brown, P., Lauder, H., and Wells, A., editors, *Education: Culture, Economy, Society*, pages 1–45. Oxford University Press, New York.

Brown, S. and Taylor, K. (2007). Religion and education: evidence from the national child development study. *Journal of Economic Behavior and Organisation*, 63(3):439–460.

Brunner, E. J. and Imazeki, J. (2005). *Tiebout choice and the voucher*. Quinnipiac University.

Bryk, A., Lee, V., and Holland, P. (1993). *Catholic schools and the common good*. Harvard University Press, Cambridge, MA.

Buddin, R. and Zimmer, R. W. (2005). *Is Charter school competition in California improving the performance of traditional public schools?* RAND Education working paper.

Burgess, S. and Briggs, A. (2006). School assignment, school choice and social mobility. *CMPO Working Paper*, 06/157.

Burgess, S., McConnell, B., Propper, C., and Wilson, D. (2007). The impact of school choice on sorting by ability and socio-economic factors in English secondary education. In Woessmann, L. and Peterson, P., editors, *Schools and the equal opportunity problem*. MIT Press, Cambridge, MA.

Burgess, S., Propper, C., Slater, H., and Wilson, D. (2005a). Who wins and who loses from school accountability? The distribution of educational gain in English secondary schools. *CMPO Working Paper*, 05/128.

Burgess, S., Propper, C., and Wilson, D. (2005b). Will more choice improve outcomes in education and health care? The evidence from economic research. CMPO.

Burgess, S. and Slater, H. (2006). Using boundary changes to estimate the impact of school competition on test scores. *CMPO Working paper*, 158.

Bush, T., Coleman, M., and Glover, D. (1993). *Managing autonomous schools: The Grant-Maintained experience*. Paul Chapman Publishing Ltd, London.

Card, D., Dooley, M., and Payne, A. A. (2007). School competition and efficiency with publicly funded Catholic schools. Mimeo. UC Berkeley.

Carr, M. and Ritter, G. (2007). Measuring the competitive effect of charter schools on student achievement in Ohio's traditional public schools. University of Arkansas.

Carrington, W. J. and Troske, K. R. (1997). On measuring segregation in samples with small units. *Journal of Business and Economic Statistics*, 15:402–409.

Census of Great Britain 1851 (1960). Religious worship in England and Wales. Report and Tables (PP1852-3), LXXXIX.

Chamberlain, T., Rutt, S., and Fletcher-Campbell, F. (2006). *Admissions: who goes where? Messages from the statistics*. National Foundation for Educational Research, Slough.

Cheshire, P. C. and Sheppard, S. C. (1995). On the price of land and the value of amenities. *Economica*, 62:247–267.

Church of England Board of Education School Committee (1972). *Crisis in church schools: A report on finance*. Church Army Press, Oxford.

Clark, D. (2007). *Politics, markets and schools: Quasi-experimental evidence on the impact of autonomy and competition from a truly revolutionary UK reform*. University of Florida.

Coldron, J. and Boulton, P. (1991). Happiness as a criterion of parental choice of school. *Journal of Education Policy*, 6:169–178.

Coldron, J., Tanner, E., Finch, S., Shipton, L., Wolstenholme, C., Willis, B., Demack, S., and Stiell, B. (2008). Secondary school admissions. *DCSF Research Report*, RR020.

Coleman, J. and Hoffer, T. (1987). *Public and private high schools*. Basic Books, New York.

Coleman, J., Hoffer, T., and Kilgore, S. (1982). *High school achievement*. Basic Books, New York.

Conway, S. (1997). The reproduction of exclusion and disadvantage: Symbolic violence and social class inequalities in 'parental choice' of secondary education. *Sociological Research Online*, 2.

Cortese, C. F., Falk, R. F., and Cohen, J. K. (1976). Further considerations on the methodological analysis of segregation indices. *American Sociological Review*, 41:630–637.

Crook, S. (1999). Ordering risks. In Lupton, D., editor, *Risk and sociocultural theory: New directions and perspectives*. Cambridge University Press, Cambridge.

Croxford, L. (2000). Is free-meal entitlement a valid measure of school intake characteristics? *Educational Research and Evaluation*, 6:317–335.

Davies, P., Adnett, N., and Mangan, J. (2002). The diversity and dynamics of competition: Evidence from two local schooling markets. *Oxford Review of Education*, 28(1):91–107.

De Fraja, G. and Landeras, P. (2006). Could do better: The effectiveness of incentives and competition in schools. *Journal of Public Economics*, 90:189–213.

Deem, R. and Wilkins, J. (1992). Governing and managing schools after ERA: The LEA experience and the GM schools alternative. In Simkins, T., Ellison, L., and Garrett, V., editors, *Implementing educational reforms: The early lessons*. Longman/BEMAS, Harlow.

Department for Children, Schools and Families (2008a). The composition of schools in England. Statistical Bulletin.

Department for Children, Schools and Families (2008b). The National Challenge: Raising standards, supporting schools. Available from http://www.dfes.gov.uk/publications/nationalchallenge/.

Department for Education (1980). *Education Reform Act*. HMSO, London.

Department for Education (1992). *Choice and Diversity: A New Framework for Schools*. HMSO, London.

Department for Education and Employment (1999). *Code of Practice on school admissions*. HMSO, London.

Department for Education and Science (1986). *City Technology Colleges: A New Choice of School*. HMSO, London.

Department for Education and Science (1988). *Education Reform Act: Local Management of Schools*. HMSO, London.

Department for Education and Skills (2001). *Schools: Achieving Success (White Paper)*. HMSO, London.

Department for Education and Skills (2003). *School admissions code of practice*, volume DfES/0031/2003. DfES, London.

Department for Education and Skills (2007a). GCSE and equivalent results in England, 2005/06. DfES Statistical First Release, Ref: 01/2007: http://www.dfes.gov.uk/rsgateway/DB/SFR/s000702/index.

Department for Education and Skills (2007b). School Admissions Code. Available from http://www/dfes.gov.uk/sacode/.

Dills, A. K. (2005). Does cream-skimming curdle the milk? A study of peer effects. *Economics of Education Review*, 24:19–28.

DiNardo, J. and Lee, D. S. (2004). Economic impacts of unionization on private sector employers: 1984-2001. *NBER Working Paper*, 10598.

Eastwood, R. and Turvey, K. (2008). Equal opportunities or loaded dice? The 2007 Admissions Code after the Brighton and Hove adjudication. *Research Intelligence*, 102.

Echols, F. and Willms, D. (1992). Alert and inert clients: The Scottish experience of parental choice of schools. *Economics of Education Review*, 11(4):339–350.

Edwards, T., Fitz, J., and Whitty, G. (1989). *The State and private education: An evaluation of the Assisted Places Scheme*. Falmer Press, London.

Epple, D. and Platt, G. J. (1998). Equilibrium and local redistribution in an urban economy when households differ in both preferences and incomes. *Journal of Urban Economics*, 43:23–51.

Epple, D. and Romano, R. E. (1998). Competition between private and public schools, vouchers and peer-group effects. *American Economic Review*, 88:33–62.

Epple, D. and Romano, R. E. (2003). Neighbourhood schools, choice, and the distribution of educational benefits. In Hoxby, C. M., editor, *The economics of school choice*, pages 227–286. University of Chicago Press, Chicago.

Evans, W. N. and Schwab, R. M. (1995). Finishing high school and starting college: Do Catholic schools make a difference? *Quarterly Journal of Economics*, 110:947–974.

Feinstein, L. and Symons, J. (1999). Attainment in secondary school. *Oxford Economic Papers*, 51:300–321.

Ferreyra, M. M. (2005). *Estimating the effects of private school vouchers in multi-district economies*. Tepper School of Business, Carnegie Mellon University.

Fiske, E. B. and Ladd, H. F. (2003). School choice in New Zealand: A cautionary tale. In Plank, D. N. and Sykes, G., editors, *Choosing choice: school choice in international perspective*. Teachers College Press, New York.

Fitz, J., Halpin, D., and Power, S. (1993a). *Grant maintained schools*. Kogan Page, London.

Fitz, J., Halpin, D., and Power, S. (1993b). *Grant-Maintained schools: Education in the market place*. Kogan Page Limited, London.

Fitz, J., Taylor, C., Gorard, S., and White, P. (2001). Local Education Authorities and the regulation of educational markets: Four case studies. *Cardiff University School of Social Sciences Occassional Paper*, 41.

Flake, G. W. (1998). *The computational beauty of nature: Computer explorations of fractals, chaos, complex systems, and adaptation*. The MIT Press, Cambridge, MA.

Flatley, J., Connolly, H., Higgins, V., Williams, J., Coldron, J., Stephenson, K., Logie, A., and Smith, N. (2001). Parents' experiences of the process of choosing a secondary school. *DfES Research Report*, 278.

Fossey, R. (1994). Open enrolment in Massachusetts: Why families choose. *Educational Evaluation and Policy Analysis*, 16:320–334.

Frölich, M. (2007). Regression discontinuity design with covariates. *IZA Discussion Paper*, 3024.

Fuchs, T. (2005). School choice under restricted mobility regimes. Paper presented at Public Choice Society Annual Meeting, Texas.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

Galindo-Rueda, F. and Vignoles, A. (2004). The heterogeneous effect of selection in secondary schools: Understanding the changing role of ability. *IZA Discussion Paper*, 1245.

Galster, G. (2002). An economic efficiency analysis of deconcentrating poverty populations. *Journal of Housing Economics*, 11(4):303–329.

Galster, G. C., Quercia, R. G., and Cortes, A. (2000). Identifying neighborhood thresholds: An empirical exploration. *Housing Policy Debate*, 11(3):701–732.

Gardner, R., Cairns, J., and Lawton, D. (2005). *Faith schools: Consensus or conflict?* RoutledgeFalmer, London.

Gay, J. D. and Greenough, J. (2000). The geographical distribution of church schools in England. Culham College Institute, Oxford.

Geller, C. R., Sjoquist, D. L., and Walker, M. B. (2001). The effect of private school competition on public school performance. *National Centre for the Study of Privatisation of Education Occasional Paper*, 15.

Gewirtz, S., Ball, S. J., and Bowe, R. (1995). *Markets, choice and equity in education*. Open University Press, Buckingham.

Gibbons, S. and Machin, S. (2003). Valuing English primary schools. *Journal of Urban Economics*, 53:197–219.

Gibbons, S., Machin, S., and Silva, O. (2008a). Choice, competition and pupil achievement. *Journal of the European Economic Association*, 6(4):912–947.

Gibbons, S. and Silva, O. (2006). Faith primary schools: Better schools or better pupils? *Centre for the Economics of Education Discussion Paper*, 72.

Gibbons, S., Silva, O., and Wilson, J. (2008b). Urban density and pupil attainment. *Journal of Urban Economics*, 63(2):631–650.

Gibbons, S. and Telhaj, S. (2006). Peer effects and pupil attainment: Evidence from secondary school transition. *Centre for the Economics of Education Discussion Paper*, 63.

Gillborn, D. (1996). Exclusions from school. Viewpoint (Number 5), Institute of Education, London.

Glanville, J. L., Sikkink, D., and Hernández, E. I. (2008). Religious involvement and educational outcomes: The role of social capital and extracurricular participation. *The Sociological Quarterly*, 49(1):105–137.

Glennerster, H. (1991). Quasi-markets and education. *Economic Journal*, 101:1268–1271.

Glennerster, H. (2002). United Kingdom education 1997-2002. *Oxford Review of Economic Policy*, 18:120–136.

Goldstein, G. S. and Pauly, M. V. (1981). Tiebout bias on the demand for local public goods. *Journal of Political Economy*, 16:131–143.

Goldstein, H. and Leckie, G. (2008). School league tables: What can they really tell us? *Significance*, 5(2):67–69.

Goldstein, H. and Noden, P. (2003). Modelling social segregation. *Oxford Review of Education*, 29:225–237.

Goldthorpe, J. H. (2000). *On Sociology: Numbers, Narratives and the Integration of Research and Theory*. Oxford University Press, Oxford.

Goodman, A. C. and Thibodeau, T. G. (1998). Housing market segmentation. *Journal of Housing Economics*, 7:121–143.

Gorard, S., Taylor, C., and Fitz, J. (2002a). Does school choice lead to 'spirals of decline'? *Journal of Education Policy*, 17:367–384.

Gorard, S., Taylor, C., and Fitz, J. (2002b). Markets in public policy: The case of the UK Education Reform Act 1988. *International Studies in Sociology of Education*, 12:23–41.

Gorard, S., Taylor, C., and Fitz, J. (2003). *Schools, markets and choice policies*. RoutledgeFalmer, London.

Grace, G. (2002). *Catholic schools: Missions, market and morality*. RoutledgeFalmer, London.

Gray, J., Goldstein, H., and Thomas, S. (2001). Predicting the future: The role of past performance in determining trends in institutional effectiveness at A level. *British Educational Research Journal*, 27:391–406.

Greene, J. P. and Forster, G. (2002). Rising to the challenge: The effect of school choice on public schools in Milwaukee and San Antonio. Manhattan Institute Civic Bulletin.

Greene, K. V. and Kang, B.-G. (2004). The effect of public and private competition on high school outputs in New York State. *Economics of Education Review*, 23:497–506.

Grogger, J. and Neal, D. (2000). Further evidence on the effects of Catholic secondary schooling. *Brookings-Wharton Papers on Urban Affairs*, 2000:151–193.

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

Hall, J. C. and Vedder, R. K. (2003). The impact of private schools on public school performance: evidence from Ohio. *Journal of Economics and Politics*, 16(1):77–92.

Halsey, A. H., Heath, A., and Ridge, J. (1980). *Origins and destinations*. Clarendon Press, Oxford.

Hastings, J. S., Kane, T. J., and Staiger, D. O. (2006). Parental preferences and school competition: Evidence from a public school choice program. *NBER Working Paper*, 11805.

Hechter, M. (1976). U.K. county data, 1851-1966 [computer file]. Colchester, Essex: UK Data Archive.

Hobbs, G. (2007). *Investigating social class inequalities in educational attainment: The effects of schools and the validity of free school meal status as a proxy for socio-economic status*. PhD thesis, Institute of Education, University of London.

Hobbs, G. and Vignoles, A. (2007). Is children's free school meal 'eligibility' a good proxy for family income? *Centre for the Economics of Education Discussion Paper*, 84.

Holmes, G. M., DeSimone, J. S., and Rupp, N. G. (2003). Does school choice increase school quality? *NBER Working Paper*, 9683.

Hotelling, H. (1929). Stability in competition. *Economic Journal*, 39:41–57.

Hoxby, C. (1994). Do private schools provide competition for public schools? *NBER Working Paper*, 4978.

Hoxby, C. (2000). Does competition among public schools benefit students and tax-payers? *American Economic Review*, 90:1209–38.

Hoxby, C. M. (2001). How school choice affects the achievement of public school students. Hoover Institution, Koret Task Force.

Hoxby, C. M. (2003a). Introduction. In Hoxby, C. M., editor, *The economics of school choice*, pages 1–22. University of Chicago Press, Chicago.

Hoxby, C. M. (2003b). School choice and school productivity: Could school choice be a tide that lifts all boats? In Hoxby, C. M., editor, *The economics of school choice*, pages 287–342. University of Chicago Press, Chicago.

Hsieh, C.-T. and Urquiola, M. (2003). *When schools compete, how do they compete? An assessment of Chile's nationwide school voucher program*. University of California, Berkeley.

Hutchens, R. M. (2004). One measure of segregation. *International Economic Review*, 45:555–578.

Iannaccone, L. R. (1997). Skewness explained, a rational choice model of religious giving. *Journal for the Scientific Study of Religion*, 36:141–157.

Jacob, B. A. and Lefgren, L. (2007). What do parents value in education? An empirical investigation of parents' revealed preferences for teachers. *Quarterly Journal of Economics*, 122(4):1603–1637.

Jepsen, C. (2003). The effectiveness of Catholic primary schooling. *Journal of Human Resources*, 38:928–941.

Johnson, H. (2006). *Reflecting on faith schools*. Routledge, Abingdon.

Jowett, S. (1995). *Allocating secondary school places: A study of policy and practice.* National Foundation for Educational Research, Slough.

Kamiya, K. (1984). On the existence and uniqueness of general equilibrium prices. *International Economic Review*, 25(1):79–83.

Katzner, D. (1989). *The Walrasian vision of the microeconomy: An elementary exposition of the structure of modern general equilibrium theory.* University of Michigan Press, MI.

Kelly, A. (2007). *School choice and student well-being: Opportunity and capability in education.* Palgrave Macmillan, Basingstoke.

Kelly, S. E. (1978). *The schools of the established church in England: A study of diocesan involvement since 1944.* PhD thesis, University of Keele.

Krupka, D. J. (2008). The stability of mixed income neighbourhoods in America. *IZA Discussion Paper*, 3370.

Ladd, H. F. and Fiske, E. B. (2003). Does competition improve teaching and learning? Evidence from New Zealand. *Education Evaluation and Policy Analysis*, 25:97–112.

Lauen, D. L. (2003). An agent based modeling approach to school choice. Mimeo. University of Chicago.

Lavy, V. (2006). From forced busing to free choice in public schools: Quasi-experimental evidence of individual and general effects. *NBER Working Paper*, 11969.

Le Grand, J. (1991). *Equity and choice.* Harper Collins, London.

Le Grand, J. (2003). *Motivation, agency, and public policy. Of knights and knaves, pawns and queens.* Oxford University Press, Oxford.

Le Grand, J. (2007). *The other invisible hand. Delivering public services through choice and competition.* Princeton University Press, Princeton, New Jersey.

Lee, D. S. (2005a). *Randomized experiments from non-random selection in U.S. house elections.* UC Berkeley.

Lee, M.-J. (2005b). *Micro-econometrics for policy, program and treatment effects.* Oxford University Press, Oxford.

Leech, D. and Campos, E. (2003). Is comprehensive education really free? A case study in the effect of secondary school admissions policies on house prices in one local area. *Journal of the Royal Statistical Society Series A*, 166:135–154.

Levačić, R. (1998). Local management of schools in England: Results after six years. *Journal of Education Policy*, 13(3):331–350.

Levačić, R. (2004). Competition and the performance of English secondary schools: further evidence. *Education Economics*, 12:179–194.

Levačić, R. and Hardman, J. (1998). Competing for resources: The impact of social disadvantage and other factors on English secondary schools' financial performance. *Oxford Review Of Education*, 24:303–328.

Levačić, R. and Hardman, J. (1999). The performance of Grant Maintained schools in England: an experiment in autonomy. *Journal of Education Policy*, 14:185–212.

Levačić, R., Jenkins, A., Vignoles, A., Steele, F., and Allen, R. (2005). Estimating the relationship between school resources and pupil attainment at key stage 3. DfES Research Report 679.

Levačić, R. and Woods, P. A. (2002). Raising school performance in the league tables (part 1): Disentangling the effects of social disadvantage. *British Educational Research Journal*, 28:207–226.

Levin, H. M. (1991). The economics of educational choice. *Economics Of Education Review*, 10(2):137–158.

Lupton, R. (2005). Social justice and school improvement: Improving the quality of schooling the poorest neighbourhoods. *British Educational Research Journal*, 31:589–604.

Machin, S. (2008). The new economics of education: Methods, evidence and policy. *Journal of Population Economics*, 21:1–19.

Mayet, G. (1996). Admissions to schools: A study of Local Education Authorities. In Glatter, R., Woods, P. A., and Bagley, C., editors, *Choice and diversity in schooling: Perspectives and prospects*, pages 166–177. Routledge, London.

McCrary, J. (2007). Manipulation of the running variables in the regression discontinuity design: A density test. *NBER Technical Working Paper*, 334.

McEwan, P. J. and Carnoy, M. (2000). The effectiveness and efficiency of private schools in Chile's voucher system. *Educational Evaluation and Policy Analysis*, 22:213–239.

Meen, G. (2005). Local housing markets and segregation in England. *Economic Outlook*, January:11–17.

Meen, G. and Meen, D. (2003). Social behaviour as a basis for modelling the urban housing market. *Urban Studies*, 40:917–935.

Ministry of Education (1944). *Education Act of 1944*. His Majesty's Stationery Office, London.

Ministry of Education (1950). *Education 1900-1950: The report of the Ministry of Education*. His Majesty's Stationery Office, London.

Neal, D. (1997). The effects of Catholic secondary schooling on educational attainment. *Journal of Labor Economics*, 15:98–123.

Nechyba, T. J. (2000). Mobility, targeting, and private-school vouchers. *American Economic Review*, 90:130–146.

Nechyba, T. J. (2003a). Introducing school choice into multidistrict public school systems. In Hoxby, C. M., editor, *The Economics of School Choice*, pages 145–194. University of Chicago Press, Chicago.

Nechyba, T. J. (2003b). School finance, spatial income segregation and the nature of communities. *Journal of Urban Economics*, 54:61–88.

Nechyba, T. J. (2004). Income and peer quality sorting in public and private schools. In Hanushek, E. A. and Welch, F., editors, *Handbook of economics of education*. North Holland, Amsterdam.

Nesheim, L. (2002). Equilibrium sorting of heterogeneous consumers across locations: Theory and empirical implications. *CeMMAP Working Paper*, 08/02.

Nguyen, A. N., Taylor, J., and Bradley, S. (2006). The estimated effect of Catholic schooling on educational outcomes using propensity score matching. *Bulletin of Economic Research*, 58(4):285–307.

Ni, Y. (2007). The impact of charter schools on the efficiency of traditional public schools: Evidence from Michigan. *National Centre for the Study of Privatisation of Education Occasional Paper*, 145.

Noden, P. (2000). Rediscovering the impact of marketisation: Dimensions of social segregation in England's secondary schools, 1994-1999. *British Journal of Sociology of Education*, 21:371–90.

Office for National Statistics (2004). *All fields postcode directory 2004 user guide v.3*. ONS Geography.

Pardey, D. (1991). *Marketing for schools*. Kogan Page, London.

Parsons, E., Chalkley, B., and Jones, A. (2000). School catchments and pupil movements: a case study in parental choice. *Educational Studies*, 26:33–48.

Pennell, H., West, A., and Hind, A. (2007). *Religious composition and admission processes of faith secondary schools in London*. Comprehensive Future, London.

Pickering, W. (1967). The 1851 religious census – a useless experiment? *British Journal of Sociology*, 18:382–407.

Pugh, G., Davies, P., and Adnett, N. (2006). Should we have faith in not-for-profit providers of schooling? *Journal of Education Policy*, 21(1):19–33.

Rapp, G. C. (2000). Agency and choice in education: Does school choice enhance the work effort of teachers? *Education Economics*, 8(1):37–63.

Reay, D. (1998). Engendering social reproduction: Mothers in the educational marketplace. *British Journal of Sociology of Education*, 19(2):195–209.

Reay, D. and Ball, S. J. (1997). 'Spoilt for choice': The working-classes and educational markets. *Oxford Review of Education*, 23:89–101.

Regnerus, M. and Elder, G. (2003). Staying on track in school: Religious influences in high- and low- risk settings. *The Journal for the Scientific Study of Religion*, 42:633–649.

Reynolds, D. (1992). School effectiveness and improvement. In Bashi, J. and Sass, Z., editors, *School effectiveness and improvement: Proceedings of the Third International Congress for School Effectiveness*. The Van Leer Jerusalem Institute, Jerusalem.

Ribchester, C. and Edwards, B. (1998). Co-operation in the countryside: small primary school clusters. *Educational Studies*, 24:281–294.

Richards, N. J. (1970). Religious controversy and the school boards 1870-1902. *British Journal of Educational Studies*, 18(2):180–196.

Robertson, D. and Symons, J. (2003). Do peer groups matter? Peer group versus schooling effects on academic attainment. *Economica*, 70:31–53.

Rosenbaum, P. and Rubin, D. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrica*, 70:41–55.

Rosenthal, L. (2003). The value of secondary schooling. *Oxford Bulletin of Economics and Statistics*, 65:329–355.

Roth, A. E. (1984). The evolution of the labor market for medican interns and residents: A case study in game theory. *Journal of Political Economy*, 92:991–1016.

Roth, A. E. (1996). Comments on Tversky's "Rational Theory and Constructive Choice". In Arrow, K. J., Colombatto, E., Perlman, M., and Schmidt, C., editors, *The rational foundations of economic behaviour*, pages 198–202. Macmillan, New York.

Rothstein, J. (2005). *Does competition among public schools benefit students and taxpayers? A comment on Hoxby (2000).* NBER working paper 11215.

Sacker, A., Schoon, I., and Bartley, M. (2002). Social inequality in educational achievement and psychosocial adjustment throughout childhood: Magnitude and mechanisms. *Social Science and Medicine*, 55:863–880.

Saltman, R. and Von Otter, C. (1992). *Planned markets and public competition.* Open University Press, Buckingham.

Sander, W. (1996). Catholic grade schools and academic achievement. *Journal of Human Resources*, 31:540–548.

Sandstrom, F. M. and Bergstrom, F. (2005). School vouchers in practice: Competition will not hurt you. *Journal of Public Economics*, 89:351–380.

Sass, T. R. (2006). Charter schools and student achievement in Florida. *American Education Finance Association*, 2006:91–122.

Sattinger, M. (1980). *Capital and the distribution of labor earnings.* North-Holland, Amsterdam.

Savage, M., Barlow, J., and et al. (1992). *Property bureaucracy and culture: Middle class formation in contemporary Britain.* Routledge, London.

Sawkins, J. W., Seaman, P. T., and Williams, H. C. (1997). Church attendance in great britain: An ordered logit approach. *Applied Economics*, 29:125–134.

Schagen, I. and Schagen, S. (2005). The impact of faith schools on pupil performance. In Gardner, R., Cairns, J., and Lawton, D., editors, *Faith schools: Consensus or conflict?*, pages 202–212. RoutledgeFalmer, London.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal Of Mathematical Sociology*, 1:143–186.

Schelling, T. C. (1978). *Micromotives and macrobehavior.* Norton, New York.

Scherer and Ross (1989). *Industrial Market Structure and Economic Performance.* Houghton Mifflin, Boston.

Schneider, M. and Buckley, J. (2002). What do parents want from schools? Evidence from the internet. *Educational Evaluation and Policy Analysis*, 24(2):133–144.

Select Committee for Education & Skills (2005). Fifth report of session 2004-05 (secondary education).

Sherratt, B. (1994). *Grant-Maintained status: Considering the options.* Longman, Harlow.

Sieg, H., Smith, V. K., Banzhof, H. S., and Walsh, R. (1999). The role of optimizing behavior in willingness-to-pay estimates for air quality. *American Journal of Agricultural Economics*, 81:1112–1117.

Simon, B. and Chitty, C. (1993). *SOS: Save our schools*. Lawrence and Wishart, London.

Simon, H. A. (1947). *Administrative behavior: A study of decision-making processes in administrative organizations*. The Free Press, New York.

Skinner, G. (2002). Religious pluralism and school provision in Britain. *Intercultural Education*, 13(2):171–181.

Skolverket (1996). Choosing schools - effects of choice in comprehensive education.

Skolverket (1999). Child care and school in figures, 1999. part 2: Children, staff, and teachers.

Stillman, A. and Maychell, K. (1986). *Choosing schools - parents, LEAs and the 1980 Education Act*. NFER-Nelson, Windsor.

Summers, A. A. and Wolfe, B. L. (1977). Do schools make a difference? *American Economic Review*, 67:639–652.

Sutton Trust (2006). The social composition of top comprehensive schools: Rates of eligibility for free school meals at the 200 highest performing comprehensive schools. Available from http://www.suttontrust.com/annualreports.asp.

Taylor, C. (2001). Hierarchies and local markets: The geography of the lived market place in secondary education provision. *Journal of Education Policy*, 16, 3:197–214.

Tesfatsion, L. S. (2006). Agent-based computational economics: A constructive approach to economic theory. In Tesfatsion, L. S. and Judd, K. L., editors, *Handbook of computational economics*, volume 2, chapter 16, pages 831–880. Elsevier, Amsterdam.

The Office of the Deputy Prime Minister (2004). *The English Indices of Deprivation 2004 (revised)*. HMSO.

Thompson, M. (1992). The experience of going Grant Maintained: The perceptions of AMMA teacher representatives. *Journal of Teacher Development*, 1:133–140.

Thrupp, M. (1999). *Schools making a difference: Let's be realistic*. Open University, Buckingham.

Tinbergen, J. (1959). On the theory of income distribution. In Klaassen, L. H., Koyck, L. M., and Witteveen, H. J., editors, *Selected papers of Jan Tinbergen*. North-Holland, Amsterdam.

Tversky, A. (1996). Rational theory and constructive choice. In Arrow, K., Colombatto, E., Perlman, M., and Schmidt, C., editors, *The rational foundations of economic behaviour*, pages 185–197. Macmillan, London.

Walford, G. (1994). *Choice and Equity in Education*. Cassell Education, London.

Walford, G. (2003). School choice in England and Wales. In Plank, D. N. and Sykes, G., editors, *Choosing choice: School choice in international perspective*. Teachers College Press, New York.

Walford, G. and Miller, H. (1991). *City Technology College*. Open University Press, Buckingham, UK.

Webber, R. and Butler, T. (2007). Classifying pupils by where they live: How well does this predict variation in their GCSE results? *Urban Studies*, 44(7):1229–1253.

West, A. and Hind, A. (2003). *Secondary school admissions in England: Exploring the extent of overt and covert selection*. Research and Information on State Education, London.

West, A., Hind, A., and Pennell, H. (2004). School admissions and 'selection' in comprehensive schools: Policy and practice. *Oxford Review of Education*, 30:347–369.

West, A. and Pennell, H. (1999). School admissions: Increasing equity, accountability and transparency. *British Journal of Educational Studies*, 46:188–200.

West, A. and Pennell, H. (2000). Publishing school examination results in England: Incentives and consequences. *Educational Studies*, 26(4):423–436.

West, A. and Pennell, H. (2002). How new is New Labour? The quasi-market and English schools 1997 to 2001. *British Journal of Educational Studies*, 50:206–224.

Wilson, D. (2003). *Which ranking? The use of alternative performance indicators in the English secondary education market*. CMPO, University of Bristol.

Wilson, D., Croxson, B., and Atkinson, A. (2006). What gets measured gets done. *Policy Studies*, 27(2):153–171.

Wilson, D. and Piebalga, A. (2008). Accurate performance measure but meaningless ranking exercise? An analysis of the English school league tables. *CMPO Working Paper*, 07/176.

Wise, A. E. (1967). *Rich schools, poor schools*. University of Chicago Press, Chicago.

Woods, P. A. (1993). Responding to the consumer: Parental choice and school effectiveness. *School Effectiveness and School Improvement*, 4:205–229.

Woods, P. A., Bagley, C., and Glatter, R. (1998). *School choice and competition*. Routledge, London.

Zoloth, B. S. (1976). Alternative measures of school segregation. *Land Economics*, 52:278–98.