

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Heterogeneous effects of High School peers on educational outcomes

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1813324> since 2022-07-04T09:56:14Z

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Heterogeneous effects of high school peers on educational outcomes

By Silvia Mendolia^a, Alfredo R. Paloyo^b, and Ian Walker^c

^a School of Accounting, Economics, and Finance, University of Wollongong, Wollongong NSW 2522, Australia, ARC Centre of Excellence for Children and Families over the Life Course (LCC), Global Labor Organization (GLO) and Institute for the Study of Labor (IZA); e-mail: smendoli@uow.edu.au

^b University of Wollongong, Leibniz-Institut für Wirtschaftsforschung (RWI), Global Labor Organization (GLO), ARC Centre of Excellence for Children and Families over the Life Course (LCC), and IZA

^c Lancaster University Management School and IZA

Abstract

We investigate peer ability effects on high-stakes test scores at ages 16 and 18, and on the probability of university attendance. To account for endogeneity in peer ability, we use the average ability of the primary school peers of one's secondary school peers, excluding those from the same primary school of the individual, as an instrumental variable for average secondary school peer ability. Our results show that average peer quality has a small effect on an individual's test scores, and a larger proportion of low-quality peers has a significantly detrimental effect on achievements of an average student. Furthermore, peer ability seems to have a stronger effect on students at the bottom of the grade distribution, especially at age 16.

JEL classifications: I20, J24

1. Introduction

Establishing the presence and size of peer effects in education is important because peer effects imply that educational interventions have multiplier effects (Glaeser *et al.*, 2003). That is, the impact of an educational intervention on an individual may propagate across a wider group of students, akin to how unvaccinated children benefit from herd immunity from vaccinated children. Moreover, the existence of heterogeneity in peer effects across the ability distribution provides a rationale for the efficient mixing of pupils in a school or in a classroom. An optimal student mix may raise the average attainment of a group in ways that other educational interventions may not be able to achieve. These considerations have implications for cost–benefit analyses of educational policy prescriptions.

However, identifying the effect of peer ability on individual achievement is complicated for reasons pertaining to the endogeneity of the peer ability measure (Angrist, 2014). In the absence of random assignment, measures of peer ability are endogenous because of self-selection into groups. For example, children attending the same school are likely to have common characteristics related to the area in which they live and their families’ socioeconomic backgrounds. The correlation between these factors and both the individual educational outcomes and the average peer characteristics causes the ordinary least-squares estimator to be biased.

Moreover, individuals also affect their own peer group as much as the peer group affects them (the so-called ‘reflection’ problem coined by Manski, 1993).¹ As a result, peer achievements are not exogenous with respect to individual educational outcomes, especially when pupils have been exposed to each other for some years. Reflection refers to the possibility that while Student A affects Student B, the latter also affects the former. If this reflection issue is not taken into account, OLS is once again rendered a biased estimator for peer effects.

The aim of the present paper is to investigate the relationship between average peer ability and individual educational attainment in high-stakes educational tests at the end of compulsory

¹ Manski (1993) distinguishes between the three channels through which individuals may have characteristics and outcomes similar to their peer group: via the endogenous effect, via exogenous effects (also called contextual effects), and via correlated effects. In our context, an endogenous effect arises if the individual’s achievement varies with the average achievement of the peer group; an exogenous effect arises if the individual’s achievement varies with the observable socioeconomic characteristics of the peer group; and correlated effects arise if the individual has similar achievements as her peers because they are subject to similar unobservable factors.

schooling at age 16, and at the end of high school at age 18, using a large, rich, and recent dataset of English teenagers. We use an identification strategy based on an instrumental variable for average peer ability to recover a credible causal estimate of peer effects and to examine whether impact heterogeneity exists across the ability distribution. The instrument is the average peer ability of the peers of a student's peers, i.e., the peers of one's peers – with the additional restriction that these 'peers of peers' have not shared an educational institution with the student.

It is the exclusion of one's own earlier peers from this peers-of-peers calculation that is the primary contribution of this paper. Failure to apply this restriction leaves open the possibility that the instrumental variable will still be correlated directly with the outcome. We believe our analysis is the first study to adopt this strategy of excluding peers of peers who once belonged to or are still among one's own peers. This allows us to overcome the reflection problem and mitigate selection bias. Since these peers are not likely to have had direct contact with a particular student in question in a schooling context, the application of this additional constraint improves upon the existing studies that use peers of peers as an instrumental variable (e.g., De Giorgi *et al.*, 2010 and Gibbons and Telhaj, 2016), which would not have been able to adequately neutralise reflection effects.

This paper locates itself within the expanding literature on peer effects in education. A number of studies find statistically significant but modest peer effects (e.g., Lefgren, 2004 and Gibbons and Telhaj, 2016) on individual test scores while others find no evidence of an effect (e.g., Lavy, Silva, and Weinhardt, 2012). The evidence is mixed for the existence of heterogeneous peer effects (see Sacerdote, 2011 for a review). Many existing papers that rely on large administrative datasets contain few relevant covariates. While containing many observations, the administrative database that forms the basis for UK studies, known as the National Pupil Database (NPD), has only a limited set of background characteristics. In contrast, our study is primarily based on the Longitudinal Study of Young People in England (LSYPE), which includes a long vector of information on the child, the family, and the school. The richness of LSYPE allows us to better control for confounding effects that threaten the conditional mean independence assumption required for causal inference. The additional covariates also allow for more precise estimates to the extent that it reduces the residual variance.

The LSYPE data allow us to identify the primary school peers of an individual's secondary school peers. This is pivotal because it forms the basis for our instrumental variable for an individual's average peer ability in high school. The maintained assumption is that these peers of peers do not have any direct effect on individuals included in the LSYPE, but rather only an indirect effect through their peers. The peers of peers must satisfy two conditions: first, they must have attended a different primary school from the student of interest; and, second, they must not be in the same high school as the student of interest.

Our results indicate that the average peer ability has no significant impact on test scores at age 16, but it has significant effects on performance at age 18 test scores, conditional on taking the test. We do not find that the average peer ability significantly affects the likelihood of attending university. The number of weak peers also has an effect on average outcomes at age 18, but not younger. However, average peer ability seems to have a stronger effect on students at the bottom of the grade distribution, especially at age 16. Finally, when we analyse the effect of the presence of weak peers on age 16 and 18 test scores across the ability distribution, we find evidence that weak peers are more detrimental to weak students than to stronger ones.

We contribute to the existing literature on peer effects in education in three main ways. First, the identification strategy is novel. This has applications to other settings where average peer characteristics are endogenous – that is, where one might be concerned about endogenous sorting and the reflection problem – but past peers of peers can be identified, such as, for example, in large linked employer–employee datasets. Second, while the existing literature based on British data analyse the impact of peers on junior high school achievement at age 14, our analysis is within the context of high-stakes educational outcomes at age 16 and at the end of high school at age 18. Third, we provide corroborative evidence for the existence of heterogeneous peer effects across the grade distribution, and we consider its consequent policy implications.

The rest of this paper is organised as follows. In Section 2, we provide a brief overview of the existing literature. We present the data and explain the peer-ability indicators and outcomes in Section 3. In Sections 4 and 5, we discuss the estimation methods and the results, respectively. Finally, we conclude in Section 6 with a discussion of policy implications.

2. Literature review

Peer effects have been estimated for a variety of outcomes, including risky health behaviours (Trogdon *et al.*, 2008; McVicar and Polanski, 2014), and a number of academic and educational outcomes (Hanushek *et al.*, 2003; Zimmerman, 2003; Carrell *et al.*, 2009; Duflo *et al.*, 2011; Lavy, Silva, and Weinhardt, 2012; Gibbons and Telhaj, 2016). Relevant to our work are studies that analyse the effect of peer ability on educational achievements in school. Many studies have looked at low-stakes outcomes for children in primary school environments and have exploited several different strategies to analyse the impact of peers in early ages (e.g., Hoxby, 2000; Hanushek *et al.*, 2003; Angrist and Lang, 2004; Lefgren, 2004; Vigdor and Nechyba, 2007; and Goux and Maurin, 2007; Ammermueller and Pischke, 2009; Landini *et al.*, 2016).

Hanushek *et al.* (2003) use data from the Texas Assessment of Academic Skills (for students in grade three through to eight) and control for fixed school, individual, and school-by-grade effects to show that peer achievements have a positive effect on individual grades, and that this effect is constant across quartiles of the grade distribution. Similarly, Lefgren (2004) uses data from Chicago public elementary schools and examine peer effects using school tracking policies. The author shows that peer effects are quite small but generally positive and significant. Angrist and Lang (2004) analyse the results of the Metropolitan Council for Educational Opportunity (METCO) program in Boston, which sends black disadvantaged students to public primary schools in high-socioeconomic-status areas; the study indicates that there is limited evidence of significant effects. A distinct strand of the literature examines peer effects in middle and secondary schools. Several studies show mostly small – but nonetheless often statistically significant – peer effects (e.g., Kang, 2007; Schindler Rangvid, 2008; Calvó-Armengol *et al.*, 2009; and Lavy, Paserman, and Schlosser, 2012).

A recent review by Sacerdote (2011) suggests that – while there is no consensus in the effect of the average peer quality on the performance of the average person – there is sufficient evidence to think that the effect is nonlinear. That is, the effect of bad (or good) peers may be different from the effect of the average, and the effect of the average may be different for bad (or good) people. In particular, Ding and Lehrer (2007) and Burke and Sass (2013) find that high-ability students benefit from

other high-ability students while Imberman *et al.* (2012) find that good peers have positive effects which are greatest for low-achieving students. Lavy, Silva, and Weinhardt, (2012) estimate a significantly negative effect from bad peers using British data, and positive effects from academically bright peers for girls.

In the UK, Bradley and Taylor (2008) estimate peer effects using information on pupils moving high schools in the last two years of their compulsory education. They show that peer effects exist and are stronger for low-ability students and non-white children. However, pupils who change schools may be systematically different from those who do not, especially when the reasons for the change can be related to school achievements. In addition, Atkinson *et al.* (2008) use a panel of school children from the southwest of England to look at the effect of the introduction of teacher-performance-related pay in England, and show significant and non-trivial peer effects while conditioning for school and teacher fixed effects.

De Giorgi *et al.* (2010) analyse peer effects in college major choices and use an identification strategy that is similar to ours. They use information on students who are enrolled in the same year in a Business faculty, and consider as peers those students who attend at least four classes together. They then select students who are not in one individual's peer group (because they do not attend classes with her/him) but are included in the group of the individual's peers (because they attend classes with some of them) and use their characteristics to instrument peers' college major choices. However, this identification strategy relies on the assumption that students who are in the same cohort – specifically, they belong in the same year of the same faculty in the same university – cannot be considered peers simply because they do not attend classes together. The implicit assumption is that non-academic interactions (e.g., common friends, social activities outside the college, etc.) do not impact on college major choice.

In Australia, McVicar *et al.* (2016) analyse peer effects on literacy outcomes in high school, using individual fixed effects (rather than instrumental variables) and separating peers who did not attend the same primary school as the individual.

The studies that are closest to ours are Lavy, Silva, and Weinhardt (2012) and Gibbons and Telhaj (2016). Both papers exploit the change in peers from primary to high school and use the NPD

to analyse the effect of peer ability measured at the end of primary school through Key Stage 2 examinations (at age 11) and on achievements in high school, measured through Key Stage 3 exams (at age 14).² Lavy, Silva, and Weinhardt (2012) use within-pupil and cross-subject regressions, and rely entirely on the variation in relative subject ability for a given group of peers. Gibbons and Telhaj (2016) exploit year-to-year changes in secondary school peer group for students changing from the same primary school to different secondary schools, and account for latent primary and secondary school fixed effects. While Lavy, Silva, and Weinhardt (2012) find no average effect of peer ability, they do find strong effects of bad peers on average outcomes. In contrast, Gibbons and Telhaj (2016) find small and significant average peer effects, but the effect does not vary significantly across the ability distribution.

This present study distinguishes itself by estimating peer effects on high-stakes outcomes at the end of high school and at entry into tertiary education. Furthermore, our identification strategy is an improvement in that we only rely on peers of peers who never had any contact with the individual. Last, both Lavy, Silva, and Weinhardt (2012) and Gibbons and Telhaj (2016) use the Pupil-Level Annual School Census (PLASC) and the NPD, which has a very limited set of family characteristics, and do not include the detailed set of socioeconomic-background variables which are available in LSYPE.³

3. Data

3.1 Institutional background

Education in England is organised into ‘Key Stages’ (KS). Children enter primary school at 4–5 years old, and move to Key Stage 1 (age 6–7). Key Stage 2 starts at age 7–8, and lasts until age 10–11 (Year 6) when children leave primary education and enter secondary school. At this point, Key Stage 3 starts (age 11–14), followed by Key Stage 4 (age 14–16). At the end of Key Stage 4, students take the General Certificate of Secondary Education exams (GCSEs), which coincides with the end of compulsory schooling. Students typically take between five to ten subjects, and passing grades in five

² Key Stages are discussed in more detail in Sec. 3.1.

³ We tested our analysis using all variables available in NPD and have found that it does not capture the effect of other important variables, such as a local deprivation index, parental education, and parental employment.

of them – including Mathematics and English – are often required to transition into academic A-levels (see next paragraph) in senior high school.

After this, students may decide to pursue further studies from age 16–18 (sometimes in a different school if their own school caters just for age 11–16). In this Key Stage 5 (age 17–18), children specialise and study more challenging subjects in preparation for their General Certificate of Education (GCE) Advanced Level examinations (the so-called ‘A-levels’). Usually, three to four subjects are studied at A-level over a two-year period, and are examined at the end of each year. Students may select subjects that depend on their academic preferences and intentions toward further education.

Local Educational Authorities (LEA) are responsible for organising the admission policies for their primary and secondary schools.⁴ Our sample includes over 640 high schools, and over 82% of them are government ‘comprehensive’ schools, while voluntary-aided and controlled schools (usually those schools with a religious denomination) – who obtain the vast majority of their funding from the public purse – provide high school education for about 15% of children.⁵

Children within a given area will often attend one of a small number of primary schools, and neighbouring children will often move on to one of several high schools. Parents are free to choose any secondary school they prefer, but when schools are oversubscribed, places are allocated according to some published criteria. Usually, looked-after children and children with special needs have priority, followed by children who have siblings in the same school, and then children living in the area, with proximity used as the tie-breaker.

In secondary schools, students are often grouped with different peers for different subjects, so they do not have a unique ‘class’ for all subjects. Students are also sometimes taught in groups of similar ability (determined after an initial observation period) for some subjects, although not all schools ‘set’ by ability, and this varies by subject, with a higher prevalence of ability setting for Mathematics

⁴ Recent policy in England (Wales, Scotland and Northern Ireland, which are not in LSYPE, have long had separate control over the shape of their education systems) have led to a more diversified population of schools with the conversion of an increasing number of publically funded comprehensive secondary schools becoming ‘Academy’ schools with greater independence from both central and local government control. Publicly-funded schools cannot select students on the basis of their ability, although some studies have suggested that schools find ways to select students on the basis of parental characteristics that might be correlated with ability (West *et al.*, 2003).

⁵ Children who have been selected by ability constitute about 2.5% of our sample. We retained them in our analysis, but their exclusion do not affect our results in any important way.

and Science and a lower incidence for English (Kutnick *et al.*, 2006). The criteria for admission exclude using ability, however, except for a very small proportion of state schools that are allowed to do this.

Some GCSE examinations (age 16), are organised in ‘tiers’, where different students sit a different test depending on their ability group, and the maximum grade that they can achieve depends on their allocated tier. As noted in Lavy, Silva, and Weinhardt (2012), a high level of ability-setting might affect measures of peer quality and might lead to downward bias in estimates of peer effects. However, our identification strategy relies on peers of peers’ ability in primary schools, where a low degree of ability setting is expected.

3.2 Dataset

The LSYPE dataset was managed by the UK Department of Education (2004–2012) and covers a wide range of topics, including family relationships, attitudes toward school, family and labour market, and some more sensitive or challenging issues, such as risky health behaviours, and personal relationships.⁶ Young people included in LSYPE were selected to be representative of all young people in England, but the survey also oversampled specific groups, in particular, young people from a low socioeconomic background. Schools and students were sampled using a two-stage sample procedure with disproportionate stratification, and students from ethnic minorities were over-sampled in order to achieve target-issued sample numbers of 1,000 in each group (Department of Education, 2011).⁷ The survey started when the adolescents were in Year 9 in 2004 at the age of 14. In the first wave of LSYPE, around 15,000 young people were interviewed across more than 700 high schools. On average, data were collected for 27 students in each school. In the first four waves, parents or guardians were also interviewed.

The records of LSYPE children can be linked to the NPD, a pupil-level administrative database of all English pupils including detailed information on pupil test scores, achievements, and school characteristics. We use this data to collect information about LSYPE children’s results in test

⁶ The LSYPE is now called ‘Next Steps’, and it is managed by the Centre for Longitudinal Studies at the University College London and funded by the Economic and Social Research Council.

⁷ We have tested our models using sample weights and main findings are unchanged.

scores at ages 11, 14, and 16, which is the minimum school-leaving age for this cohort. This occurs at the end of a stage of Key Stage 4, and culminates in the GCSEs exams. The LSYPE data includes information on Key Stage 5 exams (with detailed grades by subject studied) and the intention to participate in and actual attendance at a higher education institution.

Our final sample includes 9,213 observations of children with non-missing information on test scores at ages 11, 14 and 16, peer test scores, and other essential information on the child's family background, coming from 640 high schools and 4,126 primary schools. When we estimate the impact of peers' ability on test scores at age 17–18, the sample becomes smaller because we only include individuals who remain in education at age 17 (Wave 4 of LSYPE). The observations from our estimation sample were not significantly different from the original data in terms of relevant observable characteristics.

3.3 Outcomes

Our interest is in analysing the effect of peer ability on academic outcomes at the end of high school and on the chances that a young person will take further studies after compulsory education. We analyse peer effects on GCSE examinations score performance (age 15–16), A-level attainments (age 17–18), the likelihood of attending university, and the likelihood of admission into a 'prestigious' higher education institution. Our dependent variables include the number of subjects with 'pass' grades (A*–C) in GCSE exams; and a binary indicator for having five GCSE passes including Mathematics and English, which is the 'gold standard' that is usually required for students to follow an academic track for progression into senior high school and beyond.

We also explore the impact of peer ability on student performance in Mathematics and Science at A-levels.⁸ As noted in Mendolia and Walker (2014), the determinants of performance in a particular subject are very hard to disentangle from overall school performance. It is particularly interesting to analyse peer effects in performance in these STEM (Science, Technology, Engineering, and Math) subjects, as the UK ranking of 15-year-old pupils in Mathematics and Science in the OECD's

⁸ We group the following subjects under 'Science': Biology, Chemistry, Physics (and any combination of two of these three subjects), Environmental Science, Psychology (as a Science), Technology, Zoology, Meteorology, Engineering Science, Other Science.

Programme for International Student Assessment (PISA) tests has been consistently falling from 2000 to 2009. Furthermore, the UK has one of the lowest shares of 15-year olds intending to pursue a STEM career among the OECD countries, and particularly lags behind in women’s aspirations to study a STEM subject and engage in a STEM career (OECD, 2012).

Finally, we examine whether average peer ability has significant impacts on the conditional probabilities of attending a university and attending one of the Russell Group institutions.⁹ Of those who stay in education after age 16, 35% attend university and, within this subsample, 23% attend an institution that is part of the Russell Group of institutions. Table 1, which provides the descriptive statistics for our outcome variables, shows that more than half of the adolescents in the sample achieved five or more GCSE exams with a passing grade between A* and C, and 42% take A-level examinations two years later.

3.4 Other independent variables

We exploit the richness of the LSYPE data and estimate three versions of our model, progressively expanding the set of covariates to account for mediators and to establish the robustness of the estimated impacts. We present estimation results for a model which controls for a basic set of individual and family characteristics, including child’s sex and achievement in the KS2 test (age 11); maternal education and marital status; employment status of both parents (Wave 4 – age 17); and, additionally, the Index of Multiple Deprivation (IMD) score, which is a measure of socioeconomic status.¹⁰ Table 2 displays the descriptive statistics for a number of the independent variables included in the analysis.

⁹ The Russell Group consists of the following 24 institutions: Birmingham, Bristol, Cambridge, Cardiff, Durham, Edinburgh, Exeter, Glasgow, Imperial College, King’s College London, Leeds, Liverpool, London School of Economics, Manchester, Newcastle, Nottingham, Oxford, Queen Mary College London, Queen’s Belfast, Sheffield, Southampton, University College London, Warwick, and York. These institutions are research intensive and committed to provide an outstanding teaching and learning experience to their students (<http://russellgroup.ac.uk/>).

¹⁰ The IMD captures a wide set of characteristics of the area where the student lives, including income, employment, health and disability, education, housing, crime, and living environment, and is therefore a very important variable to deal with spatial clustering and capture families’ SES. We also estimate models which exclude the IMD score. Additionally, we have results for models with extended controls to embrace individual ethnic background and for some school characteristics, such as government region, number of students, religious denomination, and the gender mix of the school. These results are presented in the Appendix.

3.5 Peer ability

The principal explanatory variable, mean peer ability, is measured through achievements in KS3 tests (age 14) for children who attended the same high school of each LSYPE child. Our instrumental variable, the mean peers-of-peers ability, is measured through achievements in KS2 tests taken at the end of primary school at age 10–11. Our analysis is limited to children who are in LSYPE, and, consequently, we do not have a complete overview of all students in a particular primary or high school. However, using students in LSYPE allows us to access to all the available information on their families and backgrounds, which are not included in NPD but could be relevant in explaining academic outcomes.¹¹

The vast majority of high schools (around 80%) draw their students from a group of four to 23 primary schools. Table 3 shows that over 70% of children have a peers-of-peers group that contains three or more students. In our sample, the size of the peers-of-peers group varies from one to 97 children. LSYPE children have a high school peer group of, on average, 15 students (in LSYPE) who come from many different primary schools (from two to 23 primary schools).¹²

To explore effect heterogeneity, we investigate the effect of low-achieving peers in high school. In particular, we analyse whether a large fraction of ‘bad peers’ is detrimental to student learning, and we use the information on the percentage of students not achieving the basic standard (called Level 5) in KS3 Mathematics.¹³ This variable is collected at the school level, and is thus related to all the students attending a particular school and not just to pupils in the LSYPE. The majority of schools in the estimation sample have a percentage of students not achieving basic standards in Maths below 30% and, as expected, there are very few schools where more than 50% of students are in this category.

¹¹ The average peer ability measure is not calculated using all peers in a school (secondary or primary), but only the subsample in LSYPE. We discuss this issue in greater detail at the end of Section 5, and also run several sensitivity tests to verify the stability of our estimates when varying the sample size.

¹² We ran sensitivity tests in order to verify the stability of our findings when we exclude individuals who only have 1, 5, 10 or 20 peers of peers. The main results are unchanged and confirm our findings on the impact of peers’ ability (see Appendix, Table 19).

¹³ Figure A.3 in the Supplementary Appendix presents the distribution of this variable.

If there are peers that are missing from the calculation of the average peer ability, it is worthwhile to examine whether they would be missing at random or whether their disappearance is a result of some shared characteristics (Ammermueller and Pischke, 2009). However, the LSYPE sample has been selected to be representative of all young people in England and schools were stratified by deprivation status in order to ensure that schools across different levels of deprivations and different ethnic groups were equally represented (Department of Education, 2011). Therefore, it seems unlikely that peers from a particular school are missing because of some non-random characteristics, at least to the extent that they are correlated with socioeconomic status.

Measurement error may arise because we do not know exactly who the individual's peers really are. Students are not necessarily influenced by their entire school or cohort group; it could be an unknown and endogenously chosen subset of students only. If there is such measurement error in the peer population average, then there will be the same measurement error in the sample average from LSYPE. Thus, the use of the LSYPE subset is not more vulnerable to this criticism than any existing NPD-based work. The typical approach to deal with this issue is to use instrumental-variables estimation (Ammermueller and Pischke, 2009), as we have done here.

4. Estimation

We begin our analysis by estimating a linear-in-means model of peer effects:

$$Y_{ih} = \bar{A}_{ih}\alpha + \mathbf{X}'_i\boldsymbol{\gamma} + \epsilon_i, \quad (1)$$

where Y_{ih} represents a particular academic outcome for individual i who is attending high school h . We define i 's high school peers as those currently attending high school h , but they could have attended a variety of primary schools. The variable \bar{A}_{ih} is the mean ability (measured by KS3 score) for LSYPE children attending high school h excluding the individual (the 'leave-one-out' mean), and \mathbf{X}_i is a vector of child and family characteristics and a constant. The parameter of interest is α , which captures the relationship between mean peer ability, \bar{A}_{ih} , and individual achievements, Y_{ih} . Under strict exogeneity, a statistically significant coefficient implies the existence of a multiplier effect from educational interventions. However, this condition is unlikely to be met without quasi-experimentally-induced variation.

To account for the endogeneity of average peer ability in Equation (1), we use an instrumental-variable approach using peers-of-peers ability in primary school (measured as KS2 achievements – age 11) as an instrument for the mean high school peer ability (measured as KS3 achievements – age 14). We estimate the model via two-stage least squares, where the first-stage equation is

$$\bar{A}_{ih} = \bar{B}_{zg}\beta + \mathbf{X}'_i\boldsymbol{\pi} + v_i. \quad (2)$$

The mean high school peer ability \bar{A}_{ih} depends on the peers-of-peers mean performance \bar{B}_{zg} in primary school of those who attended primary schools z , different from the one attended by the individual, and currently attending high school g , where $h \neq g$. The exclusion restriction states that the ability of the high school students' peers-of-peers in primary school do not affect the high schools' student achievements directly, but only through its indirect impact via the student's current peers in high school.¹⁴

One natural concern in the estimation of this model is that selection of secondary schools on the basis of unobservable characteristics could be driving the main findings. Parents choose the school for their children (or at least the area where they live) and, thus, individuals who attend the same high school are likely to have some common background characteristics. Therefore, results could be affected by the impact of students' socioeconomic background on grades, i.e., peer estimates might include the effect of peers' background characteristics (together with peers' ability), as noted in Lavy, Silva, and Weinhardt (2012). We therefore control for a variety of parental characteristics (including marital status, labour-force status, and educational attainment) and for the IMD score in our model to mitigate this concern.

Our instrument is based on peers of peers who do not attend the same high school as the individual and did not attend the same primary school. Around 80% of high schools in the estimation sample have more than four primary 'feeder' schools, and, therefore, the peers of peers (who now at-

¹⁴ There is no reason for our measure of ability to be restricted to KS3 performance because, not only have individuals in our sample never met their peers of peers in primary school, these peers-of-peers children have also gone to different high schools. Therefore, the reflection problem does not arise. For completeness, we also present results throughout where we measure peer ability using peers' KS2 and GCSE instrumented by their peers-of-peers score at KS2 level.

tend a different high school) are likely to have come from an area with different socioeconomic characteristics. Given the considerable number of primary and high schools in the data, there is no reason to believe that these differences are systematic and peers of peers are a selected group.¹⁵

As noted in Gibbons *et al.* (2013), most households can choose between more than one school from their area of residence and, on average, students in the same cohort living in the same neighbourhood attend just one of a handful of different local secondary schools. Furthermore, a typical English secondary school is attended by pupils living in more than 60 Output Areas, the smallest proxy for neighbourhood (Gibbons *et al.*, 2013). However, neighbourhood composition has a very limited effect on test scores once one controls for family socioeconomic characteristics (Gibbons *et al.*, 2013), and the LSYPE data allow us to take into account many of these factors.

Thus far, we have assumed that peer effects are homogeneous in the sense that the relationship between peer ability and individual achievements is the same for each student in the ability distribution. However, peer effects are likely to be heterogeneous: a student's susceptibility to peer effects may vary according to one's ability. For example, peer ability might have stronger effects on weak students than on strong students, or the presence of a group of weak students might have different effects on weak students than on strong ones (see, e.g., Kang 2007).

We use quantile regression to examine the potential heterogeneous effects of peer ability at different points of the achievement distributions. We estimate the effect of the average peer ability for students at different quantiles of the GCSE and A-level score distributions. In order to deal with the endogeneity of peer ability in high school, we use IV quantile regression (Lee, 2007; Chernozhukov *et al.*, 2010; Chernozhukov *et al.*, 2015) which has been used in a similar context in Kang (2007).¹⁶

A parametric version of the estimator proposed by Lee (2007) is used in the estimation. Following Lee (2007), we estimate

¹⁵ We might expect the peers of peers to be more likely to have correlated effects with the high school peers in the data the fewer the high schools there are in a given area. For example, this might be the case with rural schools where choices may be geographically limited. It may also be the case in large high schools. We show that our headline results are robust to differentiating between rural and urban high schools, and those coming from small or large high schools (where small high schools are defined as those with less than 1,000 students, which is about the median high school size in our data).

¹⁶ The analysis is performed using the Stata routine `cqiv` with the `uncensored` option (Chernozhukov *et al.*, 2011).

$$Q_{Y_{ih}|A,X} = \bar{A}_{ih}\alpha + \mathbf{X}'_i\boldsymbol{\gamma} + \varepsilon_i, \quad (3)$$

where $Q_{Y_{ih}|A,X}$ is the quantile distribution of Y . The first-step linear quantile regression is modelled as

$$Q_{\bar{A}_{ih}|Z} = \bar{B}_{zt}\beta + \mathbf{X}'_i\boldsymbol{\pi} + \nu_j, \quad (4)$$

where $Q_{\bar{A}_{ih}|Z}$ is the quantile distribution of peer ability \bar{A}_{ih} , and β and $\boldsymbol{\pi}$ can be estimated by a quantile regression of \bar{A} on \bar{B} and \mathbf{X} .

Finally, we follow Lavy, Paserman, and Schlosser (2012) and estimate the effect of having low-ability peers in high school. We use the information on the percentage of students not achieving the basic standard (called Level 5) in KS3 (or KS2) Mathematics, which is available in the LSYPE dataset for each high school. In order to deal with the potential endogeneity of this variable, we apply the same strategy as in the previous model and instrument it with a variable indicating the percentage of students not achieving basic standards in KS2 Mathematics in peers-of-peers primary schools.

5. Results

5.1 Estimates from OLS and 2SLS

The main results are presented in Table 4, where the regressions control for individual and family characteristics, including the IMD score. Each row is a separate outcome variable; adjacent columns compares OLS and IV; the three super columns refer to different ways of measuring peer ability (i.e., peers KS2 age 10 score, their KS3 age 14 score, or the number GCSE A*–C grades at age 16).¹⁷ Note again that these are cohort-level by school-level peers and not classroom-level peers.

¹⁷ Results for less and more parsimonious specifications are provided in the Supplementary Appendix, Tables 14 and 15. It is important to show the stability of our main results when controlling for the IMD variable, as it is well-known that family socioeconomic status is a strong predictor of educational achievements later in life. The results in the Appendix show similar patterns as Table 4. We also tested our main results by including an additional indicator of economic disadvantage of the primary school attended (specifically, the percentage of students eligible for free school meal). The substantive results were unaffected. First-stage regression results from 2SLS are presented in the Appendix as well (Table 16).

Despite the long vector of control variables, the method of least squares is still likely a biased estimator for peer effects. The direction of the bias is unclear. OLS could recover overestimated coefficients because peers are endogenous, but they could also be underestimated because peer performance is measured with error. With respect to the latter, information on peers is limited to students being in the same school, but we cannot actually observe whether two pupils had direct interaction – a virtually omnipresent problem in this area. Moreover, we are measuring the average performance of peers using a subsample of individuals who are in the same high school, i.e., only those who are captured in the LSYPE dataset. Therefore, we cannot *a priori* sign the direction of bias in OLS.

The OLS results presented in Table 4 are statistically significant and suggest that improving peer ability would correlate positively with individual achievements at age 16–17 (GCSE exams) and at age 17–18 (A-level exams). The coefficient sizes are nontrivial. For example, a one-standard deviation increase in average peer KS3 score is associated with a 4.5% increase in the chances of having five or more GCSE with A*–C (the mean is 52%). To put this in context, a standard deviation in KS3 or KS2 would be approximately the difference between the score of a child of a university graduate mother and the score of one with a mother having only sub-GCSE qualification.

The IV approach produces more null results. We do not find any significant effect of peer ability on the chances of attending university or of gaining admission to an elite higher education institution. However, the OLS results for performance in test scores at age 17–18 (A-levels) carry over into the IV setting, although most of the estimated coefficients are now larger, which suggests that there could be substantial attenuation bias using OLS due to measurement error. It is also possible that IV results are higher than OLS because the estimated effect is a local average treatment effect (LATE), i.e., the effect is valid for the subgroup of individuals for whom peer effects are particularly strong.

The IV results show that a one-standard-deviation increase in the average KS3 score of the peers increases the probability of taking A-levels by about 13% (the mean is 42%); and the effect on average A-level score increases by 39.82 points, which is equivalent to 28% of a standard deviation. Peer ability also significantly increases performance in A-level Mathematics by 21.35 points (38% of

a standard deviation). These results are broadly consistent in terms of size and significance irrespective of whether we use peer KS2, KS3, or GCSE as the measure of mean peer ability.

Table 5 presents estimation results for the effect of having low-achieving peers. Each row is a different outcome, and we compare OLS and IV results across columns. We measure peer ability by KS2 and KS3 scores. The IV results indicate that being in a school with a 10% larger proportion of peers who do not achieve basic standards in Mathematics significantly decreases performance in test scores at age 14 (KS3) by about 0.20 to 0.30 points (equivalent to 3 to 5% of a standard deviation). While low-quality peers do not seem to affect GCSE performance, the results suggest that they do decrease A-level results (−14 to −18 points or 10 to 13% of a standard deviation), as well as the probability of taking A-levels in Maths and Sciences and results in A-level Maths.¹⁸ These results are similar to the corresponding effect in Lavy, Silva, and Weinhardt (2012), who show that a 10% decrease in the proportion of ‘bad’ peers at school is associated with an improvement of approximately 10 to 11% of a standard deviation of the within-pupil KS3 distribution for students.

Table 4 suggests that the average peer ability does not significantly affect performance at GCSE level (age 16–17), and it affects the performance at A-levels (age 17–18) at the 10% significance level. This is contrary to our prior expectations and we can only speculate on the reasons for this difference. One possible explanation is that the nature of A-levels differs from that for GCSEs. In particular, it seems likely that good performance at A-levels requires a greater degree of understanding and a lesser degree of rote learning than GCSE, and that peer effects act through enhancing study skills that affect understanding more than it affects memorisation.

Tables 4 and 5 also suggest that the causal effect of peer ability on university entrance is not significant, despite the apparent effect on A-level points. This is consistent with the fact that entry to university is competitive and depends on relative rather than on absolute performance at A levels.

¹⁸ We explore the sensitivity of these results to school size and urban/rural in the Appendix (see Tables A.11 and A.12), which shows that small and large schools have insignificantly different IV results (although the *F*-statistic indicates a weakness in the instrument in the case of large schools), as do urban and rural schools.

5.2 Effect heterogeneity

The lack of significance at the GCSE level could simply be due to the fact that the effect of peer ability on individual achievements is heterogeneous, i.e., that peer ability matters a lot for students at a particular point of the grade distribution, and very little for others. For example, it is possible that some students suffer (or benefit) more from their peers' ability, and, in particular, it is possible that weaker students are more heavily influenced by their peers' behaviour and achievements in class. For this reason, we follow Kang (2007) and use quantile regressions in order to analyse the potential heterogeneity of peer interactions.

Results from the estimation using quantile regression are summarised in Table 6 using the benchmark specification where average peer KS3 performance is used as a measure of peer ability. Our results confirm the main findings in the previous literature (e.g., Kang, 2007 and Carrell *et al.*, 2009), which indicate that peer effects are stronger at the bottom of the grade distribution. A one-standard-deviation increase in average peer KS3 score increases the number of GCSEs with Grades A*–C by about 0.75 for students in the 20th quantile while the effect vanishes for students in the top half of the grade distribution. Increased peer ability improves the score in A-level exams by over 50 points for students in the bottom quintile of the grade distribution, while the effect is much smaller for students at the top (see Figures A.4 and A.5 in the Supplementary Appendix for graphical evidence).

The estimated relationship between low-ability peers and outcomes confirm that an increase in the percentage of low-ability peers is associated with negative outcomes for students at the bottom of the GCSE grade distribution (see Figures A.6 and A.7 in the Supplementary Appendix). Ten percent more high school peers who do not achieve basic standards in Mathematics predicts the number of GCSEs at passing A*–C grades to decrease by about 0.3 for students in the 20th percentile of the grade distribution, while the association is significantly smaller and then vanishes for top students. A 10% increase in the proportion of peers who do not achieve basic standards in Mathematics has a negative association that ranges from 19 (10th percentile of A-level distribution) to 12 points (80th percentile) in the individual A-level score.

The potential endogeneity of peer ability is taken into account using quantile instrumental variable regression, and the results are presented in Tables A.7–A.10 (in the Supplementary Appendix). Average peer effects become generally insignificant across the GCSE distribution, but, in the case of the effect of low-quality peers, the results presented in Table A.8 suggest that the effect of low-quality peers on GCSE results is especially relevant at the bottom of the distribution. The results also confirm that peer effects on A-level performance are significant across the grade distribution, especially when we look at the impact of low-quality peers in Table A.10.

5.3 Robustness

We test our main results using three sensitivity analyses. First, we re-estimate the model comparing schools with less than 1,000 students enrolled to those with more than 1,000. Large schools will typically draw from a larger number of junior schools and are more likely to implement setting by ability groups in some subjects. For this reason, we expect peer effects to be stronger in small schools, where there are a limited number of peers with which to interact. Similarly, Lavy, Silva, and Weinhardt (2012) address the ability-tracking problem by restricting their estimation to the smallest 50% of secondary schools in England. The results are similar across different school sizes with respect to the impact of low-quality peers, and seem stronger in small schools, especially when we look at the impact on the performance in A-level in Mathematics (see Appendix Table A.11).

Second, we re-estimate results comparing high schools that are in regions that are largely rural with high schools in urban areas.¹⁹ In rural areas, school choice is likely to be limited, and the student population is more likely to be homogeneous. Peer effects from low-quality peers are very similar across the two sub-samples, but the effect from average peer quality is stronger in non-rural schools (see Supplementary Appendix Table A.12).

As a final sensitivity test, we estimate a model with primary-school fixed effects in order to take into consideration the common unobserved characteristics of children who attended the same primary school (see Supplementary Appendix Table A.13). Unfortunately, our data do not allow estimating a model with high-school fixed effects, as we only have one observation of average peer KS3

¹⁹ We adopt the definition of rural areas used in the Family Resource Survey data.

score for all children attending the same high school, and this would be perfectly collinear with the fixed effect. The results are consistent with the previous findings from the OLS and IV estimates. Interestingly, in the fixed-effects model, peer ability has a significant effect on the probability of attending university.²⁰

We also investigate whether our results could be affected by the use of LSYPE, which is a subsample of the full NPD. The NPD only includes a very limited set of family characteristics, and, in particular, it does not include the detailed set of parental socioeconomic background variables which are available in LSYPE. Adding more peers by moving to NPD would only affect the precision of the estimates – at a known rate of \sqrt{N} , where N is the number of observations. However, dropping the covariates that LSYPE provides would induce bias (Pischke, 2007). Nonetheless, we ran two sensitivity tests in order to verify the stability of our results across subsamples with different sizes.

First, we extracted some random subsamples of each high school in LSYPE, including 25, 50, and 75% of students, and we re-estimated the models on each subsample. The results from this test show an overall stability of the effect of peers' ability and a loss in precision as the sample size gets smaller. This test shows that, given that LSYPE is a random sample of NPD, the size of the estimates would not be affected by using a bigger sample, and using all peers from NPD would simply increase precision (Supplementary Appendix Table A.18).

Second, we ordered LSYPE schools by the number of peers of peers that are sampled in LSYPE, and we estimated results in different subsamples progressively excluding schools with very small numbers of peers of peers. This test shows the approximate stability of the estimated peer effects and the loss in precision as the sample size gets smaller (Supplementary Appendix Table A.19). Peer effects seem stronger when we limit the sample to individuals with lots of peers of peers: the coefficients are generally higher when we limit to greater than 10 or greater than 20 peers of peers.

²⁰ The results for other independent variables are reported in Supplementary Appendix Table A.7. Not surprisingly, individual ability (measured through the KS2 score) and family socioeconomic status (and, in particular, maternal education) are strong determinant of academic achievements.

6. Conclusion

The principal contributions are as follows. First, we analyse peer effects on high-stakes outcomes at the end of high school using a very rich and recent dataset. We also propose and apply a novel identification strategy based on the peers of peers who have had no direct interaction with the student of interest. These peers of peers are primary school peers of an individual's high school peers who attended different primary and high schools from the individual, and we use information pertaining to them as an instrument for high school average peer ability. The maintained assumption is that, since these peers of peers have never been in school with the individual, they do not have a direct effect on her or his achievements.

To investigate heterogeneous peer effects, we considered the effect of being in a school with a high proportion of low-achieving peers, and we investigated the effect of peer ability across the grade distribution. Our findings show that average peer ability has a moderate effect on performance in GCSE exams at age 16, and the effect is most relevant for students at the bottom of the grade distribution. In particular, being in a school with a high proportion of low-achieving peers is particularly detrimental for the achievements of students in the bottom quartile of the GCSE distribution.

The results for A-levels are less heterogeneous and show that increased peer quality is significantly beneficial for all students across the grade distribution. These are stable to the introduction of a more detailed set of independent variables, including individual, family and school characteristics, and robust as well to IV regression and primary-school fixed effects. Our results are broadly consistent with previous findings from the literature, in particular, with Gibbons and Telhaj (2016) and Lavy, Silva, and Weinhardt (2012). A 10% larger proportion of low quality peers decreases A-level results by about 10% of a standard deviation. These results are consistent with the corresponding effect in Lavy, Silva, and Weinhardt (2012), who show that a 10% decrease in the proportion of 'bad' peers at school is associated with an improvement of approximately 10 to 11% of a standard deviation of the within-pupil KS3 distribution for students.

There is evidence for complementarities between students of different abilities generated by this study. Even if it is particularly complex to draw clear policy implications related to students' ability mixing, we believe that these results show the detrimental effect of grouping low-ability students with peers from similar ability levels. This is particularly relevant to the potential reforms of school admission policies, which could allow schools a greater freedom to select students on the basis of ability, as well as to recent reforms that created 'Free Schools' and 'Academy Schools'. Recent media and policy debates have raised concerns regarding the possibility of narrow achievement gaps between children from poor and rich socio-economic background through the creation of new types of schools, with greater freedom on curriculum choices and admission policies. In this context, it is worth being cautious, as our findings show there exists detrimental effects of peer quality on the most vulnerable students, who already are more likely to come from disadvantaged families. We believe this is an important insight for policymakers contemplating school reforms with the aim to improve outcomes for pupils from the most disadvantaged backgrounds.

Supplementary material

The data used in the paper is confidential but can be accessed at LSYPE through the UK Data Archive at <https://discover.ukdataservice.ac.uk/series/?sn=2000030>. Supplementary material – the online appendix and the stata.do files to replicate the results – are available on the OUP website.

Acknowledgements

We are grateful to the Department for Education Longitudinal Team, especially Tim Thair and Claire Baker, for providing the LSYPE data and their assistance in understanding it. The authors are solely responsible for the views expressed here. We are grateful for comments from seminar participants at Lancaster University; Melbourne Institute of Applied Economic and Social Research; Centre for Health Economics – Monash Business School; The University of Sydney; the 2016 European Association of Labour Economics conference; SFI Copenhagen; and the University of the Philippines School of Economics.

References

- Ammermueller, A. and Pischke, J.-S. (2009). Peer effects in European primary schools: Evidence from the Progress in International Reading Literacy Study. *Journal of Labor Economics* 27, 315–348.
- Angrist, J.D. (2014). The perils of peer effects. *Labour Economics* 30, 98–108.
- Angrist, J.D. and Lang, K. (2004). Does school integration generate peer effects? Evidence from Boston's METCO Program. *The American Economic Review* 94, 1613–1634.
- Atkinson, A., Burgess, S., Gregg, P., Propper, C., and Proud, S. (2008). The impact of classroom peer groups on pupil GCSE results. CMPO Working Paper 08/197, University of Bristol.
- Bradley, S. and Taylor, J. (2008). Do peers matter? Estimation of peer effects from pupil mobility between schools. LERG Discussion Paper 13, CEPM, University of Queensland, Brisbane.
- Burke, M. and Sass, T. (2013). Classroom peer effects and student achievement. *Journal of Labor Economics* 31, 51–82.
- Chernozhukov, V., Fernandez-Val, I., and Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica* 78, 1093–1125.
- Chernozhukov, V., Fernandez-Val, I., Han, S., and Kowalski, A. (2011). CQIV: Stata module to perform censored quantile instrumental variable regression. <http://EconPapers.repec.org/RePEc:boc:bocode:s457478> (last accessed 15 January 2018).
- Chernozhukov, V., Fernandez-Val, I., and Kowalski, A. (2015). Quantile regression with censoring and endogeneity. *Journal of Econometrics* 186, 201–221.
- Calvo-Arnengol, A., Pattachini, E., and Zenou, Y. (2009). Peer effects and social networks in education. *The Review of Economic Studies* 76, 1239–1267.
- Carrell, S., Fullerton, R., and West, J. (2009). Does your cohort matter? Measuring peer effects in college achievement. *Journal of Labor Economics* 27, 439–464.
- De Giorgi, G., Pellizzari, M., and Redaelli, S. (2010) Identification of Social Interactions through Partially Overlapping Peer Groups. *American Economic Journal: Applied Economics* 2, 241–275.
- Ding, W. and Lehrer, S. (2007). Do peers affect student achievement in China's secondary schools? *The Review of Economics and Statistics* 89, 300–312.
- Duflo, E., Dupas, P. and Kremer, M. (2011). Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in Kenya. *The American Economic Review* 101, 1739–1774.

- Gibbons, S, Silva, O., and Weinhardt, F.. (2013). Everybody needs good neighbours? Evidence from students' outcomes in England. *Economic Journal* 123, 831-874.
- Gibbons, S., and Telhaj, S. (2016). Peer effects: Evidence from secondary school transition in England. *Oxford Bulletin of Economics and Statistics* 78, 548-575.
- Glaeser, E., Sacerdote, B., and Scheinkman, J. (2003). The social multiplier. *Journal of the European Economic Association* 1, 345–353.
- Goux, D., and Maurin, E. (2007). Close neighbours matter: Neighbourhood effects on early performance at school. *Economic Journal* 117, 1–24.
- Hanushek, E. A., Kain, J., Markman, J., and Rivkin, S. (2003). Does peer ability affect student achievement? *Journal of Applied Econometrics* 18, 527–544.
- Hoxby, C.M. (2000). Peer effects in the classroom: Learning from gender and race variation. NBER Working Paper No. 7867, National Bureau of Economic Research, Cambridge, MA.
- Imberman, S., Kugler, A., and Sacerdote, B. (2012). Katrina's children: Evidence on the structure of peer effects from hurricane evacuees. *The American Economic Review* 102, 2048–2082.
- Kang, C. (2007). Classroom peer effects and academic achievement: Quasi-randomization evidence from South Korea. *Journal of Urban Economics* 61, 458–495.
- Kutnick, P., Hodgekinson, S., Sebba, J., Humphreys, S., Galton, M., Steward, S., Blatchford, P., and Baines, B. (2006). Pupil grouping strategies and practices at Key Stage 2 and 3: Case studies of 24 schools in England. Research Report 796, Department for Education and Skills, London.
- Landini, F., Montinari, N., Pin, P., and Piovesan, M. (2016) Friendship network in the classroom: Parents bias on peer effects. *Journal of Economic Behaviour and Organization* 129, 56–73.
- Lavy, V., Paserman, D., and Schlosser, A. (2012) Inside the black box of ability peer effects: evidence from the variation in high and low achievers in the classroom, *Economic Journal*, 122, 208–237.
- Lavy, V., and Schlosser, A. (2011). Mechanisms and impacts of gender peer effects at school. *American Economic Journal: Applied Economics* 3, 1–33.
- Lavy, V., Silva, O., and Weinhardt, F. (2012). The good, the bad, and the average: Evidence on ability peer effects in schools. *Journal of Labor Economics* 30, 367–414.
- Lee, S. (2007). Endogeneity in quantile regression models: A control function approach. *Journal of Econometrics* 141, 1131–1158.

- Lefgren, L. (2004). Educational peer effects and the Chicago public schools. *Journal of Urban Economics* 56, 169–191.
- Manski, C.F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies* 60, 531–542.
- McVicar, D. and Polanski, A. (2014). Peer effects in UK adolescent substance use: Never mind the classmates? *Oxford Bulletin of Economics and Statistics* 76, 589–604.
- McVicar, D., Moschion, J., and Ryan, C. (2016). Achievement Effects from New Peers: Who Matters to Whom?. Melbourne Institute Working Paper Series Working Paper No. 17/16.
- Mendolia, S. and Walker, I. (2014). The effect of personality traits on subject choice and performance in high school: Evidence from an English cohort *Economics of Education Review* 43, 47–65.
- OECD (2012). *Education at a Glance*. OECD, Paris.
- Sacerdote, B. (2011). Peer effects in education: How might they work, how big are they and how much do we know thus far? *Handbook of Economics of Education* 3, 247–277.
- Schindler Rangvid, B. (2008). School composition effects in Denmark: quantile regression evidence from PISA 2000 in Dustman, C. Fitzenberger, B., and Machin, S. (eds) *The Economics of Education and Training*, Physica-Verlag, Heidelberg, 179–208
- Stock, J.H. and Yogo, M. (2005), Testing for Weak Instruments in Linear IV Regression. Ch. 5 in J.H. Stock and D.W.K. Andrews (eds), *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg*, Cambridge University Press, Cambridge
- Trogdon, J., Nonnemaker, J., and Pais, J. (2008). Peer effects in adolescent overweight. *Journal of Health Economics* 27, 1388–1399.
- Vigdor, J. and Nechyba, T. (2007). Peer effects in North Carolina public schools, in *Schools and the Equal Opportunity Problem*, in L. Woessmann and P. E. Peterson (eds), MIT Press, Cambridge, MA, 73–102.
- West, A., Hind, A. and Pennell, H. (2003). Secondary schools in London: Admissions criteria and cream skimming. London: Research and Information on State Education Trust.
<http://risetrust.org.uk/pdfs/london.pdf> (last accessed 15 January 2018).
- Zimmerman, D. (2003). Peer effects in academic outcomes: Evidence from a natural experiment. *The Review of Economics and Statistics* 85, 9–23.