Do Cuts to Higher Education Subsidies Reduce Economic Spillovers? Evidence from England's Tuition Fee Reform

Luke Walsh [1]

[1]Department of Economics, University of Liverpool

Abstract

A large literature examines the private returns to higher education, and a smaller literature attempts to quantify the positive externalities generated by graduates. Yet there is limited evidence on whether these spillovers are affected by changes in higher education subsidies. This paper studies the impact of the 2012 increase in UK university tuition fees on graduate spillovers to industry. Using administrative education and labour market data, I construct a Bartik-style shift-share instrument that combines pre-reform industry dependence on graduates from different fields of study with subject-specific changes in graduate supply induced by the tuition fee reform. The resulting variation is used to estimate the effect of graduate supply on industry-level outcomes, including patenting activity, non-graduate wages, and output. The tuition fee reform significantly reduced graduate supply to more exposed industries. I find little evidence that these reductions affected innovation, productivity, or wages. These findings suggest that reductions in higher education subsidies may generate limited indirect economic costs, at least in the medium run.

Keywords: Higher Education, Spillovers, Returns to Education

1 Introduction

Governments subsidise higher education because degrees benefit not just the individuals who hold them, but the firms, workers, and communities around them. A large empirical literature documents these spillovers: an additional percentage point in the local college-graduate share raises the wages of non-graduate workers by between 1.6 and 1.9 per cent (Moretti, 2004a), boosts plant-level productivity (Moretti, 2004b), and accelerates regional patenting and growth (Aghion et al., 2009; Valero and Van Reenen, 2019). These externalities provide the canonical justification for public subsidy: because individuals cannot capture the full social return to their degree, private enrolment decisions will be socially sub-optimal without intervention (Pigou, 1920). Yet despite this theoretical foundation, governments across the developed world have substantially reduced direct subsidies to higher education over the past three decades, shifting the cost of degrees from the taxpayer to the student (Barr, 2004). Whether this retrenchment destroys the spillovers that justified the subsidy in the first place is one of the most consequential open questions in the economics of education.

Despite the breadth of this literature, existing evidence answers a subtly different question from the one that matters for policy. The social returns literature overwhelmingly estimates the spillover effects of average increases in graduate supply such as university expansions, new campus openings, or cross-city variation in college shares. They typically find that these effects are positive, particularly for STEM-intensive or innovation-complementary human capital (Winters, 2014; Iranzo and Peri, 2009; Toivanen and Väänänen, 2016). What is almost entirely absent is evidence on the spillovers generated by the specific students affected by subsidy cuts. Reducing public subsidies does not change the average graduate; it changes the marginal one. Under an income-contingent loan system, higher fees raise the effective cost of a degree only for students whose expected lifetime earnings fall short of full loan repayment, precisely those in lower-return subjects and fields (Chapman, 2006). The students deterred by subsidy cuts are therefore systematically different from the average graduate whose spillovers the existing literature measures, and there is no reason to expect the two effects to coincide. Indeed, the evidence suggests spillovers are highly heterogeneous: STEM and innovation-intensive graduates generate substantially larger externalities than non-STEM graduates (Winters, 2014), and even within STEM, policy-induced changes in the composition of entrants can reduce aggregate invention by reallocating talent away from innovation-intensive careers (Bianchi and Giorcelli, 2020). Whether subsidy reductions destroy meaningful spillovers therefore depends critically on which students are at the margin, a question the existing literature is not designed to answer.

This paper asks directly whether reducing public subsidies to higher education destroys the spillovers that justify them. I exploit the 2012 tripling of the undergraduate tuition fee cap in England from £3,375 to £9,000 as a large, plausibly exogenous shock to the effective public subsidy per student. I estimate its effect on four outcomes that the literature identifies as the principal channels through which graduate human capital generates external benefits: patents filed, firm R&D expenditure, non-graduate wages, and aggregate industry output. Crucially, I do not ask whether graduates generate spillovers on average, but whether the specific reduction in graduate supply induced by a subsidy cut propagates into measurable losses in innovation, productivity, and wages. This distinction matters because the 2012 reform operated through a well-defined margin: under the income-contingent loan system introduced in 2006, the fee increase raised the expected cost of a degree only for students in subjects where lifetime earnings fall short of full loan repayment, generating differential supply contractions across subject groups that were determined by the pre-reform earnings structure of each field rather than by post-reform industry conditions. I exploit this cross-subject heterogeneity using a Bartik-style shift-share design that interacts the subject-specific supply shock with the pre-reform dependence of each industry on graduates from each subject, producing plausibly exogenous variation in the graduate intensity of industries that is driven entirely by which subjects each industry happened to recruit from before the reform.

Our analysis draws on three primary data sources. The backbone of the empirical strategy is the Longitudinal Education Outcomes (LEO) dataset, an administrative panel that links individual HESA student records to HMRC tax records, allowing us to track England-domiciled graduates from university enrolment through to their first post-degree employment spell. From LEO we construct a subject × industry × year panel of graduate counts, log graduate counts, and graduate shares, covering all graduates entering the labour market between 2010 and 2019 across twenty-two CAH subject groups and 2-digit SIC industries. The Bartik instrument is constructed by interacting pre-reform industry exposure weights with subject-level log changes in graduate supply relative to the same pre-reform baseline, so that industries which historically recruited heavily from subjects that subsequently contracted receive greater instrument values. Spillover outcomes are measured at the industry × year level using four sources: EPO patent applications assigned to UK firms from Dimas et al. (2025); business R&D expenditure from the ONS Business Enterprise Research and Development (BERD) survey; gross value added from the ONS national accounts; and average wages of non-graduate workers aged 40–60 constructed from all available waves of *Understanding Society* (University of Essex, ISER, 2023). All specifications include industry fixed effects and a linear time trend.[1], with standard errors clustered at the industry level.

For three of our four outcomes—patents filed, gross value added, and non-graduate wages—we find no evidence that the reduction in graduate supply induced by the 2012 tuition fee reform generated measurable spillover losses at the industry level; the estimates for output and innovation are precise zeros, tight enough to rule out spillovers of the magnitude implied by leading estimates of human-capital externalities. The one exception, a negative association with industrial R&D, we attribute to industry composition rather than a causal effect of graduate supply and set aside; we return to it in the appendix. First-stage estimates confirm that the Bartik instrument is a strong predictor of industry-level graduate intensity, with industries more exposed to subject-specific supply contractions experiencing significantly lower graduate counts in the post-reform period, validating the relevance of the identification strategy. Despite this, reduced-form estimates of the instrument's direct effect on each spillover outcome are small, precisely estimated, and statistically indistinguishable from zero. Two-stage least squares estimates are similarly null, with confidence intervals tight enough to rule out economically meaningful effects of the kind documented in the existing spillovers literature. These results are robust across alternative definitions of graduate supply, different sample periods, and varying levels of industry aggregation. The null findings should not be interpreted as evidence that graduates generate no spillovers on average but rather that the margin of students deterred by subsidy cuts does not appear to be the margin that drives aggregate spillover effects. Graduates induced to exit by higher fees are disproportionately concentrated in lower-return subjects whose contribution to industry-level innovation, productivity, and wages is too diffuse to detect at the level of aggregation studied.

This paper speaks to three strands of the existing literature. The first and largest is the social returns to education literature, which documents positive spillovers from graduate human capital on wages, productivity, and innovation. Seminal contributions by Moretti (2004a) and Moretti (2004b) establish

[1] Results are robust to replacing the linear time trend with year fixed effects. The linear trend is included to capture the secular increase in graduate supply observed prior to the tuition fee reform.

that increases in the local college-graduate share raise the wages of non-graduate workers and boost plant-level productivity in the United States, while Iranzo and Peri (2009) show that an additional year of college education per worker raises state-level total factor productivity by between five and six per cent. The effects are heterogeneous: Winters (2014) documents that STEM graduates generate substantially larger wage externalities than non-STEM graduates, and Leten et al. (2014) find that graduate supply and university research generate distinct and complementary effects on firm-level patenting in Italian manufacturing. These average spillover estimates, while compelling, identify the effect of increases in graduate supply rather than the consequences of policy-induced contractions at the margin. The second strand examines spillovers from higher education policy shocks more directly. Aghion et al. (2009) use political shocks to state education appropriations to show that research and four-year college spending raise regional growth and patenting, particularly in states at the technological frontier. Valero and Van Reenen (2019) document positive effects of university expansion on regional GDP per capita across 78 countries, while Schlegel et al. (2022) find that the establishment of universities of applied sciences raises average firm profits by almost 20 per cent in treated Swiss municipalities. Closer to our setting, Bianchi and Giorcelli (2020) show that a policy reform expanding STEM access in Italy altered the occupational sorting of high-ability students in ways that reduced aggregate inventorship, highlighting that the composition of policy-induced entrants matters as much as their quantity. The third strand is the shift-share literature in labour and regional economics. Following Bartik (1991) and the identification framework formalised by Goldsmith-Pinkham et al. (2020), shift-share designs have become a standard tool for generating plausibly exogenous variation in local labour market conditions. Our application of this framework to subject-industry exposure weights is, to our knowledge, novel in the higher education spillovers literature.

This paper makes three contributions. First, we provide the first direct evidence on whether reductions in public subsidies to higher education destroy the spillovers that provide their economic rationale. Prior work has established that graduates generate positive externalities on average, and a separate literature has documented the effects of fee reforms on students' own enrolment and labour market outcomes (Belfield et al., 2017; Britton et al., 2019) [THESE CITATIONS ARE PLACEHOLDERS, ??]. What has been missing is evidence connecting subsidy reductions to the indirect effects on the workers, firms, and industries around them. Our null results fill this gap and carry a direct policy implication: governments can meaningfully reduce higher education subsidies without generating measurable losses in innovation, productivity, or non-graduate wages, at least over the medium run and at the margin of students affected by the English fee reform. Second, we introduce a novel identification strategy to the higher education spillovers literature by combining administrative graduate-level microdata with a Bartik shift-share design that exploits cross-subject heterogeneity in the enrolment response to a fee shock. This approach isolates variation in industry-level graduate intensity that is driven by pre-determined subject exposure rather than post-reform demand conditions, and is applicable to other settings where policy shocks generate differential supply contractions across fields. Third, we contribute to the growing literature on the heterogeneity of human capital spillovers by showing that the identity of the marginal student matters for whether spillover losses are detectable. The graduates most sensitive to fee increases are concentrated in subjects with lower expected returns, and our results suggest that it is precisely this group whose absence leaves no measurable trace in industry-level outcomes.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on the history of higher education funding in England, documents the differential impact of the 2012 fee reform across subject groups, and traces how this subject-level heterogeneity translates into variation in industry-level graduate exposure. Section 3 describes the data sources and the construction of the graduate supply measures, the Bartik instrument, and the spillover outcome variables. Section 4 sets out the empirical strategy, presenting the shift-share design and the first-stage, reduced-form, and two-stage least squares estimating equations. Section 5 presents the main results, reports first-stage diagnostics, and discusses the robustness of the null findings to alternative specifications. 6 discusses the checks used to validate the empirical results presented in the prior section. Section 7 concludes and discusses the implications of the findings for the fiscal case for public university funding.

2 Background

The funding of higher education in England has undergone three substantial reforms since the late 1990s, each of which shifted a greater share of the cost from the state to the individual student. Prior to 1998, undergraduate tuition was free at the point of entry for domestic students, with the government funding universities directly through block grants administered by the Higher Education Funding Council for England (HEFCE). The Teaching and Higher Education Act 1998 introduced upfront annual fees of

£1,000, subject to means-tested waivers for low-income students, marking the first explicit transfer of tuition costs to students (House of Commons Library, 2017).

The more consequential structural change came with the Higher Education Act 2004, which took effect from the 2006/07 academic year. This reform abolished upfront fee payments and replaced them with deferred, income-contingent loans repayable only once a graduate's earnings exceeded a specified threshold. The fee cap was simultaneously raised to £3,000 per year. The introduction of income-contingent loans was significant because it severed the link between the upfront cost of higher education and access to it for credit-constrained individuals. Under the pre-2006 system, prospective students from low-income households faced a genuine liquidity barrier: fees were due at enrolment, and capital market imperfections meant that some individuals could not borrow against their future earnings to finance them (Chapman, 2006; Barr, 2004). The income-contingent structure resolved this by ensuring that repayment was contingent on realised post-graduation earnings, effectively insuring students against downside labour market risk and eliminating the need for collateral or upfront payment (Chapman, 2006).

A consequence of this design, which is central to the identification strategy in this paper, is that from 2006 onwards the decision to attend university is no longer materially affected by credit constraints. Fee increases under the income-contingent system therefore operate through a single channel: the expected lifetime return to the degree net of expected loan repayments. Students who anticipate sufficient earnings to repay their loan in full face a higher effective cost; those who do not—because their lifetime earnings will fall below the repayment threshold—are largely insulated from fee increases. The 2006 reform therefore transforms fee variation into a signal about expected returns rather than a liquidity shock.

The third and most dramatic reform came in 2012/13, when the Coalition Government raised the maximum annual fee cap from £3,375 to £9,000 following the independent Browne Review (Browne, 2010). In practice, the vast majority of English universities moved immediately to the £9,000 ceiling (House of Commons Library, 2017). Critically, this increase in fee income to institutions was explicitly designed to substitute for direct public funding rather than supplement it. The HEFCE teaching grant for predominantly classroom-taught subjects was removed entirely, with HEFCE allocating teaching funding only to high-cost laboratory and clinical subjects from 2012/13 onwards (London Economics, 2014; HEFCE, 2013). The consequence was a rapid and sustained collapse in direct state support for university teaching: funding council grants for teaching fell by 78 per cent in real terms between 2010/11 and 2021/22 (House of Commons Library, 2023). Expressed as a share of institutional income, funding body grants declined from 39 per cent in 2005/06 to around 11 per cent by the mid-2020s, while fee income rose to account for more than half of total university revenues (House of Commons Library, 2024). The 2012 reform was therefore not simply a price increase for students—it was a deliberate withdrawal of public subsidy from higher education, with fee income acting as the replacement funding mechanism.

2.1 The Aggregate Stagnation of Graduate Supply

The tripling of tuition fees in 2012 coincided with a marked stagnation in the total number of graduates entering the English labour market. Figure 1 plots the total number of England-domiciled graduates from English higher education institutions over the sample period. The figure documents a pronounced flattening of graduate supply following the reform cohort, in stark contrast to the sustained growth observed throughout the pre-reform period. While aggregate enrolment did not collapse, the rate of growth fell dramatically, with the post-reform trend lying substantially below any reasonable counterfactual extrapolation of the pre-reform trajectory. This pattern is consistent with the reform operating through the returns channel described in Section 2: prospective students on the margin of attendance, for whom the expected lifetime return to a degree was insufficient to justify the increased debt burden, chose not to enrol. The aggregate effect of these marginal decisions was a persistent shortfall in graduate supply relative to trend.

2.2 Differential Effects Across Subject Groups

The aggregate stagnation in graduate supply masks substantial heterogeneity across subject groups. Because the income-contingent loan system ties the effective cost of a degree to expected lifetime earnings, the 2012 reform increased the perceived cost most sharply for subjects with lower or more uncertain graduate premia. Students in high-wage, high-certainty fields, where lifetime returns comfortably exceed the repayment threshold, faced a relatively small change in expected net cost, while students in fields with lower average earnings or greater earnings uncertainty faced a more substantial increase in their expected debt burden. The reform therefore generated differential supply shocks across subjects, with the largest contractions concentrated in fields where the graduate wage premium is lowest.

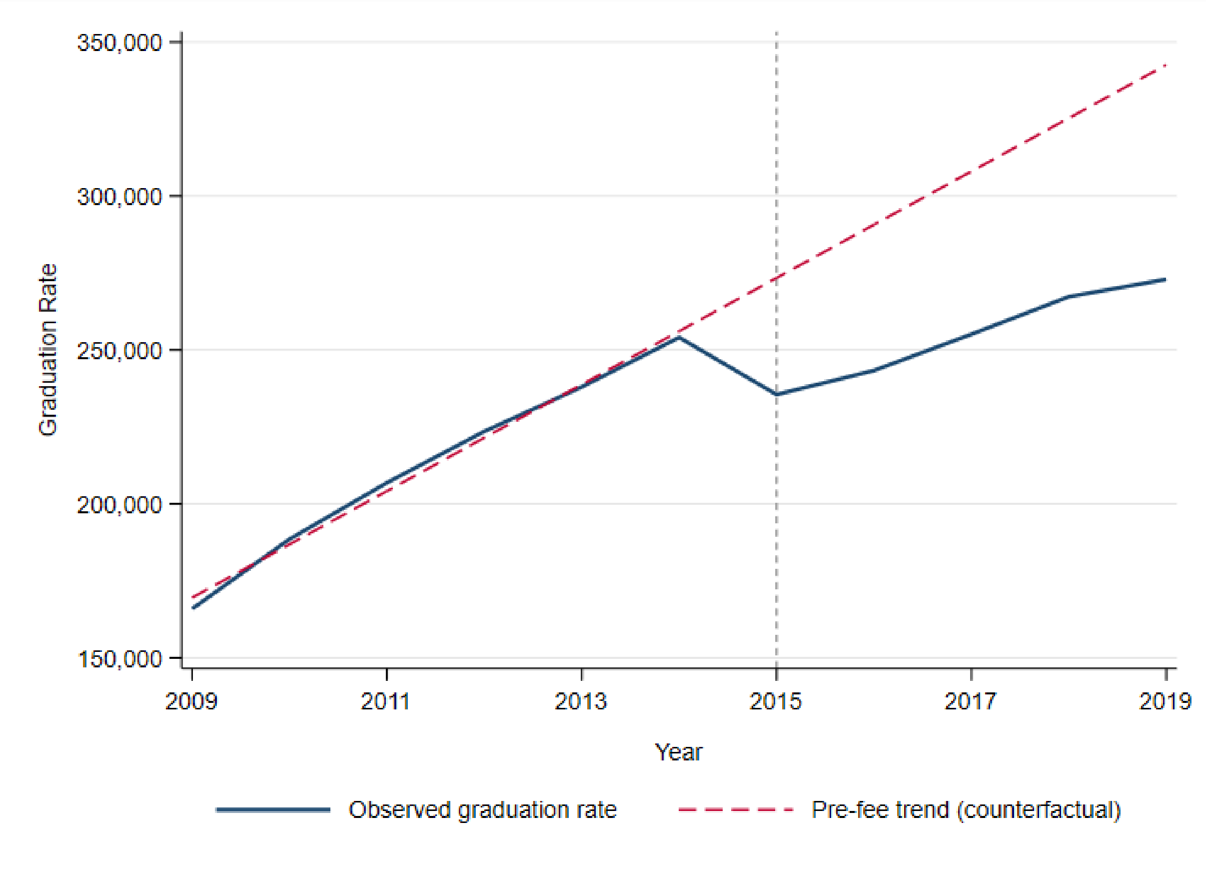

Figure 1: Total England-domiciled graduates graduating from university, 2009–2019. The dashed vertical line indicates the first cohort affected by the 2012 tuition fee reform. The dashed horizontal line indicates the pre-reform average graduate count over 2010–2014.

Figure 2 plots graduate numbers for four illustrative JACS subject areas over the sample period. In each panel the solid line is the observed graduate count and the dashed line is a linear pre-reform trend, fitted over 2010–2014 and extrapolated forward as a counterfactual; the vertical line marks 2014, the last cohort to enrol before the reform took effect. Because the first reform-affected entrants graduate from 2015 onwards, any divergence between the observed series and its counterfactual should emerge only to the right of this line—as it does in every panel. The four areas span the full range of post-reform trajectories.

The first pattern, *Subjects Allied to Medicine* (top left), is one of no discernible response: the observed series remains indistinguishable from its extrapolated trend throughout, continuing its pre-reform ascent without interruption. This is consistent with these subjects producing graduates with sufficiently high and certain earnings that the increased fee burden had negligible effect at the enrolment margin. A second pattern, illustrated by *Engineering* (top right), is one of temporary stagnation followed by recovery: graduate numbers flatten for roughly two years after 2014 before resuming growth, recovering strongly through the end of the sample although remaining somewhat below the extrapolated counterfactual. For a high-return field, this is consistent with short-run uncertainty about returns depressing entry before resolving over time. A third pattern, seen in *Education* (bottom left), is one of persistent stagnation: the series plateaus shortly after 2014 and remains broadly flat, opening a steadily widening gap below the counterfactual without any sharp break. This is consistent with a modest but durable deterrent effect at the enrolment margin. Finally, *Mass Communications* (bottom right) exhibits an outright collapse: having risen with trend to a peak around 2014, graduate numbers fall abruptly the following year and remain depressed thereafter, far below the counterfactual and with no sign of recovery. This is the sharpest and most persistent contraction in the figure, and sits naturally with the low and uncertain returns characteristic of the field.

These patterns are the *shifts* in the Bartik instrument described in Section 4. What matters for identification is not the level of graduate supply in any subject, but the change in that supply relative to the pre-reform baseline. Crucially, this change is driven by the aggregate fee shock rather than by industry-specific demand conditions. The cross-subject variation in Figure 2 is the key source of heterogeneity that allows the shift-share design to separately identify the effect of the reform on different industries, depending on their pre-reform reliance on each subject group.

2.3 Transmission to Industry-Level Graduate Exposure

The differential subject-level supply shocks described above translate into industry-level variation in graduate exposure through the pre-reform composition of each industry's graduate workforce. Industries that happened to draw disproportionately on subject groups that experienced the largest post-reform contractions received a larger negative shock to their graduate intake, while industries concentrated in unaffected subject groups were largely insulated. This is precisely the shift-share logic formalised in equation (4): the instrument aggregates subject-level supply shocks using fixed pre-reform industry weights, so that industry-level exposure is determined entirely by the historical composition of the workforce rather than by post-reform demand conditions.

Table 1 illustrates this transmission mechanism for Human Health Activities (UK SIC 2007 division 86), a sector whose pre-reform composition maps particularly clearly onto the supply shocks documented above. The table reports the share of the sector's graduate intake drawn from each JACS subject area, averaged over the pre-reform period (2010–2014).The intake is highly concentrated: Subjects Allied to Medicine and Biological sciences alone account for some 61% of the sector's graduates, while its reliance on most other areas is much lower. Second, and decisive for the magnitude of the shock the sector inherits, this concentration falls on subjects whose supply was largely insulated from the reform. The sector's single largest dependence—Subjects Allied to Medicine, at 34.7%—is precisely the area shown in Figure 2 to have continued on its pre-reform trajectory without interruption. Conversely, the areas that stagnated or collapsed after the reform—Mass Communications, Education, and Engineering—together account for under 5% of its graduate intake.

The shift-share logic therefore predicts that Human Health Activities should experience only a *small* negative exposure shock: the subjects on which it depends held up, and the subjects whose supply contracted barely feature in its workforce. The mechanism is most visible in the counterfactual. Had an otherwise identical industry instead drawn the bulk of its graduates from Mass Communications or Education, the same fixed pre-reform weights would have transmitted those sharp and persistent contractions into a correspondingly large decline in its graduate intake. It is this interaction between fixed composition and differential supply shocks that the instrument exploits, and it is what allows two industries facing the same aggregate reform to experience very different exposure.

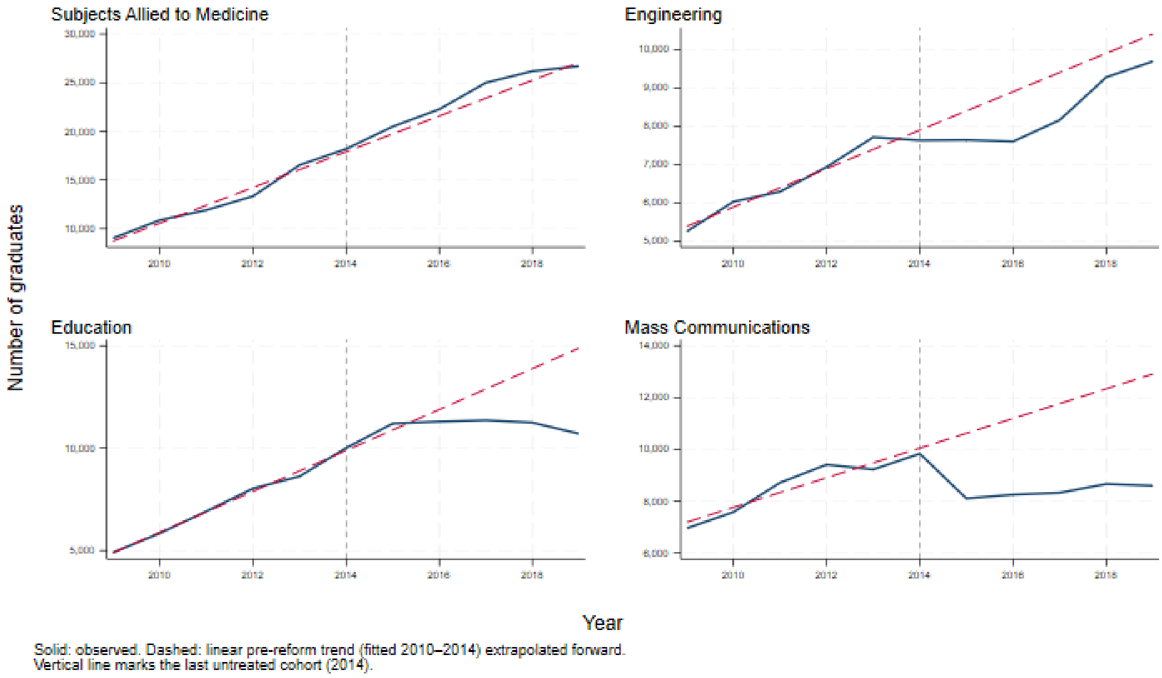

Figure 2: Graduate numbers for four selected JACS subject areas. In each panel the solid line is the observed graduate count and the dashed line is a linear pre-reform trend, fitted over 2010–2014 and extrapolated forward; the vertical line marks 2014, the last cohort to enrol before the 2012 tuition fee reform took effect (the first affected cohort graduates from 2015). The panels illustrate four distinct post-reform trajectories: no response (Subjects Allied to Medicine); temporary stagnation followed by recovery (Engineering); persistent stagnation below trend (Education); and outright collapse (Mass Communications).

Table 1: Graduate Intake Composition of Human Health Activities (SIC 86)

Subject area	Share (%)
Subjects allied to medicine	34.7
Biological sciences	26.0
Social studies	8.1
Business & administration	6.9
Creative arts & design	4.9
Linguistics & classics	3.8
Physical sciences	3.0
History & philosophy	2.9
Law	2.7
Mass communications	2.0
Education	1.5
Engineering	0.8
European languages	0.7
Architecture & planning	0.5
Veterinary science & agriculture	0.4
Combined	0.3
Technologies	0.3
Medicine & dentistry	0.2
Other languages	0.2
Total	100.0

Notes: Each row reports the share of graduate hires in the Human Health Activities sector (UK SIC 2007 division 86) who studied in the given JACS subject area, computed over the pre-reform period (2010–2014). Shares sum to 100 across subject areas. The two largest areas account for 61% of the sector's graduate intake, and the Herfindahl index of subject concentration is 0.21. Columns may not sum exactly to 100 due to rounding.

This variation across industries in their pre-reform subject composition is what generates the identifying variation in the empirical strategy. Industries did not choose their subject composition in response to the 2012 reform; these weights were determined by decades of occupational and educational sorting that predates the policy change. They are therefore plausibly orthogonal to the post-reform industry-specific shocks that might independently affect the spillover outcomes of interest, satisfying the key identifying assumption of the shift-share design (Goldsmith-Pinkham et al., 2020).

3 Data

The primary data source is the Longitudinal Education Outcomes (LEO) dataset, an administrative panel compiled by the Department for Education that links individual records across school, higher education, and employment. LEO combines the Higher Education Statistics Agency (HESA) student record with HM Revenue and Customs (HMRC) Pay As You Earn (PAYE) tax records, enabling graduates to be tracked from university enrolment through to labour market entry without the attrition problems that affect survey-based alternatives.

The analysis sample comprises all individuals domiciled in England who graduated from a UK higher education institution between 2010 and 2019. Restricting to England-domiciled students ensures comparability with the 2012 tuition fee reform, under which the fee cap for English students at English institutions rose from £3,375 to £9,000 per year. Students domiciled in Scotland, Wales, and Northern Ireland faced distinct fee regimes and are therefore excluded.

For each graduate, the outcome of interest is their first sustained employment spell observed outside of higher education following degree completion.[2] Industries are defined at the 2-digit Standard Industrial Classification (SIC) level, and subjects are defined using the Common Aggregation Hierarchy (CAH) top-level letter groupings—for example, Group A (Medicine and Dentistry), Group B (Subjects Allied to

[2]The treatment of graduates who proceed directly to postgraduate study before entering the labour market is left to be addressed as a robustness consideration; results are unlikely to be sensitive to this given that postgraduate study is itself concentrated in subjects and cohorts that differ systematically from the marginal students most affected by the fee reform.

Medicine), Group C (Biological and Sport Sciences), and so on—yielding twenty-two broad disciplinary categories. The primary analytical dataset is structured at the subject × industry × year level.

3.1 First-Stage Outcome Variable

Let n_{sit} denote the number of England-domiciled graduates from subject s observed in their first post-degree employment spell in industry i in year t. The outcome is its natural logarithm, $\ln n_{sit}$, which allows the first-stage estimates to be interpreted in proportional terms.

3.2 Bartik Instrument Construction

The identification strategy exploits the 2012 tuition fee increase as a source of plausibly exogenous variation in graduate supply. The instrument takes a Bartik shift-share form (Bartik, 1991; Goldsmith-Pinkham et al., 2020), interacting a measure of the subject-specific supply shock with a measure of each industry's pre-reform exposure to that subject.

The *shift* is defined as the growth rate in total graduate output from each CAH subject group between the post-reform period and the pre-reform baseline. The pre-reform baseline for subject s is the average annual graduate count over 2010–2014, prior to the cohorts most directly affected by the 2012 reform completing their degrees,

$$\bar{n}_s^{\text{pre}} = \frac{1}{5} \sum_{t=2010}^{2014} \sum_i n_{sit}. \tag{1}$$

The shift therefore captures the degree to which each subject experienced a contraction in graduate supply following the reform, relative to its own pre-reform trend.

The *share* measures the initial reliance of each industry on graduates from each subject, computed as the average over the pre-reform period of the fraction of an industry's graduates drawn from subject s,

$$\omega_{si}^{\text{pre}} = \frac{1}{5} \sum_{t=2010}^{2014} \frac{n_{sit}}{\sum_{s'} n_{s'it}}. \tag{2}$$

The Bartik instrument for industry i in year t is then

$$B_{it} = \sum_s \omega_{si}^{\text{pre}} \cdot g_{st}, \tag{3}$$

where g_{st} is the growth rate of total graduates from subject s by year t relative to the pre-reform baseline $\bar{n}_s^{\text{pre}}$. Industries that happened to draw disproportionately on subjects which subsequently experienced the largest contractions in graduate supply receive greater instrument values, generating variation in graduate intensity that is driven by the interaction of pre-determined exposure weights with an aggregate supply shock rather than by industry-level demand conditions.

3.3 Outcome Data

Four outcome variables are constructed at the industry × year level and merged onto the LEO-derived panel. Each captures a distinct channel through which graduate human capital may generate spillovers for the broader economy.

Industry-level Gross Value Added (GVA) is taken from the Office for National Statistics (ONS) national accounts, which publishes annual current-price GVA by 2-digit Standard Industrial Classification (SIC) division. These series provide a comprehensive measure of aggregate productive output within each industry and serve as the broadest test of whether changes in graduate supply translate into economy-wide productivity effects.

Data on business R&D expenditure come from the ONS Business Enterprise Research and Development (BERD) survey, which records annual R&D spending by UK businesses at the 2-digit SIC level. R&D expenditure is a standard empirical proxy for knowledge accumulation and innovation investment, and is a natural outcome variable in the human capital spillovers literature given the role graduates are hypothesised to play in generating and absorbing new knowledge within firms (see, e.g., Moretti, 2004b).

Patent counts are drawn from Dimas et al. (2025), which provides a longitudinal dataset of EPO patent applications and grants at the 2-digit ISIC Rev.4/NACE Rev.2 sector level for 24 European countries from 1985 to 2020, with patents assigned to sectors using weights derived from sector-technology field correspondence tables applied to OECD EPO data. The UK subset of this dataset is extracted and matched to 2-digit SIC industries using the standard NACE Rev.2 to SIC concordance.[3]. Annual counts of patent applications are aggregated to the industry × year level to form the outcome variable.

Data on non-graduate wages are drawn from *Understanding Society*, the UK Household Longitudinal Study, which has been conducted annually since approximately 2009 and covers all available waves through to the present. Non-graduates are defined as individuals with no higher education qualification. To isolate workers whose wages are most plausibly affected by graduate spillovers rather than direct human capital accumulation, the sample is restricted to those aged 40–60, an age range by which most career-relevant sorting into industries has already occurred and in which the confound of own educational upgrading is minimal. Average log hourly wages for this group are constructed at the 2-digit SIC × year level and constitute the fourth outcome variable. This measure captures the extent to which an influx of graduates raises the productivity, and hence earnings, of co-workers who themselves did not attend university, the core spillover channel emphasised in the local labour markets literature (Moretti, 2004b; Acemoglu et al., 2004).

4 Empirical Strategy

The central empirical challenge is that the number of graduates entering an industry in any given year is endogenous to conditions in that industry. Firms in expanding, high-wage sectors attract more graduates, meaning that a naive regression of spillover outcomes on graduate supply conflates the causal effect of graduates with the response of graduates to favourable industry conditions. To address this, I construct a Bartik-style shift-share instrument (Bartik, 1991; Goldsmith-Pinkham et al., 2020) that generates plausibly exogenous variation in industry-level graduate intensity by interacting pre-determined industry exposure weights with an aggregate supply shock driven by the 2012 tuition fee reform.

The identifying variation is as follows. The 2012 reform raised the tuition fee cap for English students from £3,375 to £9,000 per year, substantially increasing the private cost of higher education. The enrolment response to this shock varied systematically across subjects, with some disciplines experiencing larger contractions in graduate supply than others. Crucially, this cross-subject variation in the supply shock was largely determined by pre-existing differences in the sensitivity of prospective students to cost across subjects—factors that are plausibly unrelated to subsequent industry-level productivity trends. Industries that happened, for historical reasons, to draw disproportionately on subjects that experienced the largest supply contractions therefore received a differential negative shock to their graduate intake. This interaction of pre-determined exposure with an aggregate shock is the source of identification.

4.1 Bartik Shift-Share Design

Let s index CAH subject groups, i index 2-digit SIC industries, and t index years. The Bartik instrument assigns to each industry an exposure-weighted average of subject-level graduate supply growth rates,

$$B_{it} = \sum_{s} \omega_{si}^{\text{pre}} \cdot g_{st}, \tag{4}$$

where ω_{si}^{pre} is the pre-reform share of industry i's graduates drawn from subject s, defined as the average over 2010–2014 of the ratio of subject-s graduates to total graduates in industry i,

$$\omega_{si}^{\text{pre}} = \frac{1}{5} \sum_{t=2010}^{2014} \frac{n_{sit}}{\sum_{s'} n_{s'it}}, \tag{5}$$

and g_{st} is the log change in total graduate output from subject s in year t relative to the pre-reform baseline,

$$g_{st} = \ln\left(\sum_{i} n_{sit}\right) - \ln\left(\bar{n}_s^{\text{pre}}\right), \qquad \bar{n}_s^{\text{pre}} = \frac{1}{5} \sum_{t=2010}^{2014} \sum_{i} n_{sit}. \tag{6}$$

By construction, B_{it} is larger in magnitude for industries whose pre-reform graduate intake was concentrated in subjects that experienced stronger growth in graduate output relative to their pre-reform baseline. The exposure weights ω_{si}^{pre} are fixed in the pre-reform period and are therefore orthogonal to post-reform

[3]NACE Rev.2 and 2-digit SIC 2007 share a common structure for the two digit mapping, so no crosswalk or aggregation is necessary here.

industry-level shocks, which is the key condition required for the instrument to be valid (Goldsmith-Pinkham et al., 2020).

The validity of the shift-share design rests on two conditions. First, the relevance condition requires that B_{it} predicts actual graduate supply in industry i, which is verified empirically via the first stage. Second, the exclusion restriction requires that the instrument affects the outcome variables only through its effect on graduate supply, and not through any independent effect on industry conditions. The main threat to exclusion would be if the 2012 reform affected aggregate consumer spending or demand in ways that differentially hit industries in proportion to their pre-reform graduate intensity. The outcome variables considered here—patents, R&D expenditure, GVA, and non-graduate wages—are less sensitive to short-run demand fluctuations than employment or output in consumer-facing sectors, which mitigates this concern. Following Goldsmith-Pinkham et al. (2020), the identifying assumption can alternatively be stated in terms of the exposure shares: the pre-reform shares ω_{si}^{pre} must be uncorrelated with unobserved determinants of industry outcomes after conditioning on controls. Since these shares reflect historical patterns of graduate sorting across industries that predate the reform, this assumption is plausible.

4.2 Estimating Equations

The empirical strategy proceeds in three steps: a first stage, a reduced form, and a two-stage least squares (2SLS) second stage. All specifications are estimated on a panel of 2-digit SIC industries observed annually from 2010 to 2019, with standard errors clustered at the industry level to account for serial correlation within industries over time.

4.2.1 First Stage

The first stage estimates the effect of the Bartik instrument on industry-level graduate supply. The estimating equation is

$$y_{it}^{\text{grads}} = \alpha + \beta B_{it} + \gamma_i + \delta_t + \varepsilon_{it}, \tag{7}$$

where y_{it}^{grads} is our measure of graduate supply in industry i and year t: the natural logarithm of the raw graduate count, σ_{it}. γ_i denotes industry fixed effects, which absorb time-invariant differences in graduate intensity across sectors, and δ_t denotes a linear time trend, which absorbs trends in graduate supply common to all industries. The coefficient of interest is β, which identifies the extent to which variation in the instrument translates into variation in actual graduate supply. A valid instrument requires β to be statistically significant with the expected sign, and the first-stage F-statistic to comfortably exceed conventional weak-instrument thresholds (Staiger and Stock, 1997).

4.2.2 Reduced Form

The reduced form estimates the intent-to-treat effect of the instrument directly on each spillover outcome, bypassing the endogenous graduate supply variable. The estimating equation is

$$y_{it}^{\text{outcome}} = \alpha + \pi B_{it} + \gamma_i + \delta_t + \varepsilon_{it}, \tag{8}$$

where y_{it}^{outcome} is one of the four spillover outcomes: log GVA, log R&D expenditure, log patent counts, or log average non-graduate wages. The reduced form coefficient π captures the total effect of a unit increase in the instrument on the outcome. A null reduced form—in which $\hat{\pi} \approx 0$ and is precisely estimated—provides direct evidence that the tuition fee shock generated no discernible spillover effects through the graduate supply channel, independently of the first-stage relationship.

4.2.3 Second Stage

The 2SLS second stage uses the Bartik instrument to recover the local average treatment effect of graduate supply on each spillover outcome. The estimating equation is

$$y_{it}^{\text{outcome}} = \alpha + \lambda \hat{y}_{it}^{\text{grads}} + \gamma_i + \delta_t + u_{it}, \tag{9}$$

where $\hat{y}_{it}^{\text{grads}}$ is the predicted value of graduate supply from the first stage in equation (7). The coefficient λ is the 2SLS estimate of the causal effect of a unit increase in graduate supply on the spillover outcome, identified off the variation in graduate intensity induced by the reform. Under the null, $\hat{\lambda} \approx 0$ across all four outcomes, indicating that the reduction in graduate supply caused by the fee reform had no measurable effect on industry-level innovation, productivity, or non-graduate wages. The 2SLS estimate is interpretable as the effect on industries that were induced to receive fewer graduates as a consequence of the reform—that is, a local average treatment effect for complier industries.

5 Results

This section asks whether the reform-induced contraction in graduate supply propagated beyond the graduates themselves to the industries that employ them. The motivating hypothesis is one of human-capital spillovers: if graduates raise the productivity of those they work alongside, then an industry that loses graduate labour should see knock-on effects on the pay of its non-graduate workers, on its output, and on its innovative activity. We test this by relating each industry's exposure to the subject-specific supply shocks to three outcomes: non-graduate pay, gross value added, and patenting. Because the exposure measure isolates variation in graduate supply driven by the aggregate fee reform rather than by industry-specific demand, a relationship between exposure and these outcomes would be difficult to attribute to anything other than the change in graduate labour. The central finding, previewed here, is that no such relationship is present: across all three outcomes the estimated spillover is a precisely estimated zero.

We begin with the first stage, which confirms that the instrument captures meaningful variation in industry graduate supply. Regressing log industry graduates on the shift-share exposure measure yields a strong and consistent relationship: the first-stage coefficient is positive and tightly estimated across every sample, and the Kleibergen–Paap rk Wald F statistic ranges from roughly 22 to 25, comfortably above the Stock–Yogo critical value of 16.38 for ten per cent maximal size. The instrument is therefore not weak, and the second-stage estimates that follow can be read as well-identified elasticities rather than artefacts of a tenuous first stage. This matters especially for the interpretation of null results: a precise zero in the second stage reflects a genuine absence of spillovers, not an inability of the instrument to move the endogenous regressor.

Table 2 presents the main estimates. Across all three outcomes, the effect of graduate supply is economically small and statistically indistinguishable from zero. The 2SLS elasticities in Panel B are −0.050 for non-graduate pay, 0.066 for gross value added, and 0.134 for patenting, none significant at conventional levels; the reduced-form estimates in Panel A carry the same message, with coefficients on the exposure measure that are likewise small and insignificant. The agreement between the two panels is reassuring: the null is unlikely to be an artefact of the instrumenting step but is already visible in the direct relationship between exposure and outcomes. The industries most exposed to the reform's graduate-supply contraction fared no differently, on any of these three margins, than industries that were largely insulated.

The value of these estimates lies in their precision, which allows us to distinguish a genuine absence of spillovers from mere statistical noise. For gross value added the 95% confidence interval on the elasticity is [−0.14, 0.27], and for patenting [−0.09, 0.35]; both rule out spillover effects of the magnitude implied by leading estimates of human-capital externalities (**?**). In economic terms, a ten per cent reduction in an industry's graduate workforce is estimated to change its output by less than three per cent in either direction, and its patenting by less than four. These are tight bounds, and they place the burden of proof on the spillover hypothesis rather than leaving the question open. The non-graduate pay estimate is the least precise of the three. The confidence interval for log of pay, [−0.55, 0.45], cannot exclude moderately sized wage spillovers so the evidence against pay externalities is weaker than that against output or innovation effects. The leave-one-out specification in Table 3, which severs any mechanical link between the instrument and the outcome, returns the same three nulls, confirming that the result is not an artefact of instrument construction.

The one outcome for which exposure is significantly related to the outcome is industrial R&D, reported separately in Appendix Table 4. We set this result aside rather than treat it as evidence of spillovers, for three reasons. First, the inference is fragile: R&D is available only at the coarsest industry aggregation used in the analysis, leaving just 58 clusters, and the series additionally suffers non-random measurement loss due to ONS disclosure control, so the cluster-robust standard errors are least trustworthy precisely where this result is estimated. Second, the sign is wrong for a spillover interpretation: the estimate is *negative*, implying that a fall in graduate supply *raises* industrial R&D, which is the opposite of what human-capital complementarity would predict. Third, the result is consistent with a composition artefact rather than a causal effect. The subjects whose supply was most resilient to the reform (subjects allied to medicine) feed disproportionately into health and service sectors that are structurally low-R&D, so industries that appear "less exposed" are also, mechanically, industries that do little R&D to begin with.

Table 2: Effect of Industry Graduate Supply on Economic Outcomes

	(1) Pay	(2) GVA	(3) Patents
Panel A. Reduced form			
Bartik exposure	−0.113	0.145	0.300
	(0.563)	(0.228)	(0.211)
Panel B. 2SLS			
ln(graduates)	−0.050	0.066	0.134
	(0.250)	(0.102)	(0.110)
Kleibergen–Paap F	23.9	22.6	23.5
Observations	673	673	655
Industries	78	77	75
Industry FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: The table reports the effect of an industry's graduate supply on three economic outcomes, all in logs. Panel A reports reduced-form estimates regressing each outcome on the shift-share (Bartik) exposure measure, which combines each industry's pre-reform subject composition with subject-level changes in graduate supply following the 2012 tuition fee reform. Panel B reports 2SLS estimates instrumenting log industry graduates with the same exposure measure; the coefficient is the elasticity of the outcome with respect to graduate supply. The Kleibergen–Paap F is the first-stage cluster-robust rk Wald statistic. All specifications include industry and year fixed effects, with standard errors clustered by industry in parentheses. Pay and patents are at the SIC two-digit division level, while GVA at coarser output-based groupings.. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Because the identifying variation in exposure is cross-industry, the industry and year fixed effects do not absorb this sorting. We therefore read the R&D estimate as reflecting which industries happen to draw on resilient subjects, not as a graduate-supply spillover, and do not include it among the main results.

What might explain the absence of detectable spillovers? The most plausible account lies in *who* the reform removed from the graduate pool. By design, a change in the price of higher education operates at the enrolment margin: the students deterred by higher fees are those whose decision to attend was closest to indifference, not the inframarginal students who would have enrolled under any plausible fee regime. This selection matters for spillovers. If human-capital externalities are generated disproportionately by the most able or most specialised graduates then a policy that thins the margin will have little effect on the stock of spillover-generating human capital, even where it visibly reduces graduate numbers. The marginal graduate foregone may simply have less to contribute to the productivity of those around her than the average graduate the spillovers literature typically has in mind. On this reading our nulls are not evidence that graduate externalities are absent in general, but that the particular slice of graduates the reform removed was not the slice through which such externalities, if they exist, predominantly flow.

The UK model of universal student loans that protect credit-constrained students from up-front costs while shifting the burden of repayment onto graduates themselves has a well-documented private logic: it preserves access while recovering costs from those who benefit. Earlier work establishes that this arrangement leaves the private returns to a degree largely intact for the individuals who enrol.[4] The present results extend the ledger to the social side. If the graduates deterred at the margin generated substantial positive externalities, then the private case for the policy would have to be weighed against an uncosted social loss of fewer spillovers to non-graduates, lower output, and less innovation. We find no evidence of such a loss. The reform's contraction in graduate supply, concentrated as it was on marginal entrants, does not appear to have imposed measurable costs on the industries that employ graduates. Taken together with the private-returns evidence, this suggests that a financing model which protects access through universal loans while placing private costs on graduates can reduce publicly-subsidised enrolment without generating the broader economic costs that a spillovers-based rationale for subsidy would predict.

[4] [reference my other paper]

6 Robustness

A potential concern with any shift-share design is that the instrument and the outcome are mechanically linked through their construction. In our setting, the subject-level supply shocks that form the "shift" component are built from national graduate totals by subject, while the outcome is measured at the industry level; because each industry's own graduates contribute to the national subject totals, an industry that is a large employer of a given subject mechanically influences the very shock it is then assigned. If sizeable, this overlap could induce a spurious relationship between exposure and outcomes, biasing the estimates toward a mechanical correlation rather than a causal effect.

Table 3: Effect of Industry Graduate Supply on Economic Outcomes: Leave-One-Out Instrument

	(1) Pay	(2) GVA	(3) Patents
Panel A. Reduced form			
Bartik exposure (LOO)	−0.060	0.166	0.300
	(0.544)	(0.221)	(0.208)
Panel B. 2SLS			
ln(graduates)	−0.030	0.085	0.147
	(0.268)	(0.108)	(0.122)
Kleibergen–Paap F	19.1	17.9	18.9
Observations	673	673	655
Industries	78	77	75
Industry FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: The table reproduces Table 2 using the leave-one-out exposure instrument, which constructs each subject-level supply shock excluding the own industry's graduates, breaking any mechanical link between the instrument and the outcome. All outcomes are in logs. Panel A reports reduced-form estimates on the leave-one-out Bartik exposure measure; Panel B reports 2SLS estimates instrumenting log industry graduates with the same measure. The Kleibergen–Paap F is the first-stage cluster-robust rk Wald statistic. All specifications include industry and year fixed effects, with standard errors clustered by industry in parentheses. Pay and patents are at the SIC two-digit division level, while GVA at coarser output-based groupings. $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

To address this we re-estimate all specifications using a leave-one-out instrument, in which each industry's own graduates are removed from the national subject aggregate before the shock is computed, so that the shift assigned to an industry is constructed entirely from graduates employed elsewhere. This severs the mechanical channel by construction while preserving the economic content of the instrument: the leave-one-out shocks remain strong predictors of industry graduate supply, with first-stage F statistics of roughly 18 to 19, still comfortably above conventional weak-instrument thresholds.

Table 3 reports the results. The estimates are materially unchanged. The 2SLS elasticities remain small and statistically insignificant across all three outcomes—−0.030 for non-graduate pay, 0.085 for gross value added, and 0.147 for patenting—and the reduced-form coefficients tell the same story. That the nulls survive the leave-one-out correction is reassuring on two counts: it confirms that the absence of measured spillovers is not an artefact of mechanical overlap between instrument and outcome, and it indicates that whatever small mechanical bias the baseline specification might contain is not what is generating the zeros. The central finding—that the reform's contraction in graduate supply produced no detectable spillovers to the industries employing graduates—is robust to this alternative construction of the instrument.

7 Conclusion

The standard economic justification for subsidising higher education is that graduates confer benefits on others such as raising the productivity of those they work alongside, contributing to innovation, and lifting wages beyond their own. If these spillovers are real and large, then a policy that reduces graduate numbers imposes a social cost over and above any private one, and the case for public subsidy rests on internalising precisely that externality. This paper has asked whether such a cost materialised when the United Kingdom sharply raised the private price of a degree through the 2012 tuition fee reform. The question is of more than historical interest: spillovers are the textbook rationale for the subsidy, so a reform that withdraws the subsidy puts that rationale directly to the test.

Our empirical strategy exploited the fact that the reform reduced graduate supply unevenly across fields of study, contracting some subjects far more than others. Using a shift-share design that interacts each industry's pre-reform reliance on particular subjects with the subsequent subject-specific contractions in graduate supply, we constructed a measure of industry exposure to the reform that is driven by the aggregate fee shock rather than by industry-specific demand. The first stage confirmed that this exposure strongly predicts industry graduate supply. The second stage, however, returned a consistent null: across non-graduate pay, gross value added, and patenting, industries more exposed to the reform's graduate-supply contraction fared no differently from those that were insulated. For output and innovation these are precisely estimated zeros, tight enough to rule out spillovers of the magnitude implied by leading estimates of human-capital externalities; for non-graduate pay the evidence is weaker but still shows no detectable effect. The nulls are robust to a leave-one-out construction of the instrument that removes any mechanical link between exposure and outcomes.

We have argued that the most plausible reading of these nulls lies in the nature of the population the reform affected. A change in the price of education operates at the enrolment margin, deterring the students closest to indifference rather than the inframarginal core who would attend under any plausible regime. If the externalities that justify subsidy are generated disproportionately by the most able or specialised graduates, then thinning the margin removes graduates who had relatively little to contribute to the productivity of those around them; the absence of measured spillovers therefore follows naturally. On this interpretation our results do not show that human-capital externalities are absent in general, only that the slice of graduates the reform removed was not the slice through which such externalities predominantly flow.

This finding speaks directly to the design of higher education finance, and it complements the private-returns evidence established in earlier work.[5] That work showed that an income-contingent system that protects credit-constrained students from up-front costs while recovering the cost of education from graduates themselves can shift the private burden of higher education without harming the individuals who enrol. The present results extend that assessment to the social side of the ledger. Had the deterred marginal graduates generated substantial positive externalities, the private case for the reform would have to be set against an uncosted social loss; we find no evidence of one. Together, the two findings suggest that a financing model which safeguards access through universal loans while placing private costs on graduates can reduce publicly-subsidised enrolment without incurring the broader economic costs that a spillovers-based rationale for subsidy would predict.

These conclusions come with limits that also mark out where further work is most needed. Our design identifies the effect of the particular supply shock the reform delivered, concentrated on marginal entrants; it is silent on what a reform reaching inframarginal or higher-ability students might do, and the marginal-student mechanism we advance, while consistent with the evidence, is an interpretation rather than a directly tested channel, since we do not observe the ability or specialisation of the deterred students. The pay estimate, less precise than the others, leaves open the possibility of moderate wage spillovers that our data cannot detect. The analysis speaks to the medium-run horizon observable within our sample window; externalities that accrue over longer periods, or through channels such as the slow accumulation of innovative capacity, may lie beyond its reach. What we can say is that, on the margins where the reform acted and over the horizon we observe, the withdrawal of subsidy did not erode the spillovers that the subsidy was meant to protect.

Acknowledgements

References

Daron Acemoglu, David H. Autor, and David Lyle. Women, war, and wages: The effect of female labor supply on the wage structure at midcentury. *Journal of Political Economy*, 112(3):497–551, 2004.

Philippe Aghion, Leah Boustan, Caroline Hoxby, and Jerome Vandenbussche. The causal impact of education on economic growth: Evidence from the United States. *Brookings Papers on Economic Activity*, 2009(1):1–73, 2009.

Nicholas Barr. *The Economics of the Welfare State*. Oxford University Press, Oxford, 4th edition, 2004.

[5][jmp – mention findings on grads/non grads].

Timothy J. Bartik. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute for Employment Research, Kalamazoo, MI, 1991.

Chris Belfield, Jack Britton, Lorraine Dearden, and Laura van der Erve. Higher education funding in England: Past, present and options for the future. *IFS Briefing Note*, BN211, 2017.

Nicola Bianchi and Michela Giorcelli. Scientific education and innovation: From technical diplomas to university STEM degrees. *Journal of the European Economic Association*, 18(5):2608–2646, 2020.

Jack Britton, Lorraine Dearden, Neil Shephard, and Anna Vignoles. How english domiciled graduate earnings vary with gender, institution attended, subject and socio-economic background. *IFS Working Paper*, W19/06, 2019.

John Browne. Securing a sustainable future for higher education: An independent review of higher education funding and student finance, 2010.

Bruce Chapman. Income contingent loans for higher education: International reforms. In Eric A. Hanushek and Finis Welch, editors, *Handbook of the Economics of Education*, volume 2, pages 1435–1503. Elsevier, Amsterdam, 2006.

P. Dimas, D. Stamopoulos, and A. Protogerou. A longitudinal dataset of sector-level patent data for Europe. *Data in Brief*, 60:112095, 2025. doi: 10.1016/j.dib.2025.112095.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.

HEFCE. Higher education funding for 2013-14. Technical Report 2013/02, Higher Education Funding Council for England, 2013.

House of Commons Library. Higher education tuition fees in England, 2017.

House of Commons Library. Higher education funding in England, 2023.

House of Commons Library. Higher education finances and funding in England, 2024.

Susana Iranzo and Giovanni Peri. Schooling externalities, technology, and productivity: Theory and evidence from U.S. states. *Review of Economics and Statistics*, 91(2):420–431, 2009.

Bart Leten, Paolo Landoni, and Bart Van Looy. Science or graduates: How do firms benefit from the proximity of universities? *Research Policy*, 43(8):1398–1412, 2014.

London Economics. The higher education fees and funding reforms in England: What is the break-even point?, 2014.

Enrico Moretti. Estimating the social return to higher education: Evidence from longitudinal and repeated cross-sectional data. *Journal of Econometrics*, 121(1–2):175–212, 2004a.

Enrico Moretti. Workers' education, spillovers, and productivity: Evidence from plant-level production functions. *American Economic Review*, 94(3):656–690, 2004b.

Arthur C. Pigou. *The Economics of Welfare*. Macmillan, London, 1920.

Tobias Schlegel, Christian Pfister, and Uschi Backes-Gellner. Tertiary education expansion and regional firm development. *Labour Economics*, 78:102207, 2022.

Douglas Staiger and James H. Stock. Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557–586, 1997.

Otto Toivanen and Lotta Väänänen. Education and invention. *Review of Economics and Statistics*, 98(2): 382–396, 2016.

University of Essex, ISER. Understanding society: UK household longitudinal study, waves 1–13, 2009–2022, 2023. SN: 6614.

Anna Valero and John Van Reenen. The economic impact of universities: Evidence from across the globe. *Economics of Education Review*, 68:53–67, 2019.

John V. Winters. STEM graduates, human capital externalities, and wages in the U.S. *Regional Science and Urban Economics*, 48:190–198, 2014.

A Appendix

Table 4: Effect of Graduate Supply on Industry R&D

	(1) Baseline	(2) Leave-one-out
Panel A. Reduced form		
Bartik exposure	-2.764^{**}	
	(1.070)	
Bartik exposure (LOO)		-2.760^{**}
		(1.040)
Panel B. 2SLS		
ln(graduates)	-1.134^{**}	-1.256^{***}
	(0.430)	(0.462)
Kleibergen–Paap F	24.9	19.9
Observations	497	497
Industries	58	58
Industry FE	Yes	Yes
Year FE	Yes	Yes

Notes: R&D is reported separately from the main results. The estimate is identified off the coarsest industry aggregation in the analysis (58 clusters), and the R&D series carries non-random measurement loss due to ONS suppression for statistical disclosure, so it is less reliable than the main-table outcomes; the negative sign is also consistent with exposure proxying for low-R&D-intensity industries rather than a causal effect of graduate supply. Column (1) uses the baseline instrument; column (2) the leave-one-out instrument, which constructs each subject-level supply shock excluding the own industry's graduates. Both panels report estimates in logs. The Kleibergen–Paap F is the first-stage cluster-robust rk Wald statistic. All specifications include industry and year fixed effects, with standard errors clustered by industry in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.