



Working Papers

www.cesifo.org/wp

Testing Means-Tested Aid

Richard Murphy
Gill Wyness

CESIFO WORKING PAPER NO. 6105
CATEGORY 5: ECONOMICS OF EDUCATION
SEPTEMBER 2016

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

Testing Means-Tested Aid

Abstract

Inequalities do not end once students enter higher education. Yet, the majority of papers on the effectiveness of education aid examine its impact on college enrolment. In this paper, we provide evidence on the causal impact of means-tested financial aid on the outcomes of students who have already enrolled in college. To do so, we exploit a unique non-salient financial aid program which varies both across and within institutions, and for which eligibility is a highly non-linear function of parental income. Our findings suggest that increasing aid by £1,000 increases the chances of gaining a good degree by 3.7 percentage points.

JEL-Codes: I220, I230, I280.

Keywords: higher education, financial aid, degree completion.

Richard Murphy
Department of Economics
University of Texas at Austin
2225 Speedway, BRB 1.116, C3100
USA - 78712 Austin, Texas
richard.murphy@austin.utexas.edu

Gill Wyness
UCL Institute of Education
University College London
20 Bedford Way
UK – London, WC1H 0AL
g.wyness@ucl.ac.uk

September 2016

We thank Damon Clark, Susan Dynarski, Caroline Hoxby, Sarah Turner and John Van Reenen for their comments and suggestions. We would also like to thank participants of the AEFP 2016, NBER Education 2015, CESifo 2015 and Royal Economic Society 2014 conferences, the Russell Sage Foundation and the CEP, LSE, University of Texas at Austin and UCL Education workshops. We also thank Steve Barnes, Charley Greenwood, Prasanthi Ramakrishnan and Claudia Hupkau for their excellent assistance in collecting and cleaning the data. Wyness acknowledges ESRC funding (ES/L000105/1). All errors are our own.

1. Introduction

Given the substantial economic returns to higher education (Leslie and Brinkman, 1998; Psacharopoulos, 1994, Blundell et al, 2000), increasing the probability of degree acquisition among disadvantaged students is important for governments' human capital accumulation and social mobility strategies. It is no surprise, then, that billions of dollars per year is spent on student aid in higher education (HE) systems around the world (OECD, 2016).

To date, the majority of student aid research has focused on its effects on the extensive margin, specifically focusing on matriculation decisions (Kane 1995, Dynarski, 2000; 2003, Seftor and Turner 2002, Nielsen et al., 2010, Dearden et al, 2014). The consensus from these studies is that aid programmes increase enrolment to the tune of around 1-3 percentage points per \$1,000.

However, there is comparatively little research estimating the causal effect of aid on the intensive margin. The perennial issue is that the vast majority of financial aid programs will impact on both the extensive and intensive margins simultaneously, making it hard to isolate the impact of financial aid on student graduation. Moreover, while much of this research has looked at the impact of merit based incentives on the outcomes of enrolled students (Scott-Clayton, 2011; Garibaldi et al., 2012, and Joensen,2013), generally finding that these incentives improves student outcomes, only a small number of studies (Bettinger, 2004; Denning, 2016) have examined the effectiveness of non-incentive based Pell Grants on the outcomes of enrolled students. However, even the Pell Grant has federal standards requiring the student to maintain a minimum GPA and credit ratio, with over 40 percent of first year community college students failing to meet them (Schudde & Scott-Clayton, 2014), meaning that even this funding is not completely unconditional.¹

¹ Pell Grants have ongoing requirements called "Satisfactory Academic Progress" standards. The federal government requires that students on federal aid maintain a certain GPA and credit ratio, as well as complete their program within a specified timeframe, in addition to reapplying for aid every year. Colleges have some flexibility in how they implement the policy, but most require a 2.0 cumulative GPA and that students complete 2/3 of their credits attempted at the end of year 1.

In the absence of incentives, why would an unconditional cash transfer have any impact on student outcomes? Traditionally, student aid is thought of as a means to reduce liquidity constraints which may limit students' ability to learn. Hence a cash transfer might affect outcomes through: a) enabling students to afford additional learning materials b) enabling them to fund a better living environment, or c) reducing the need to work during college, meaning they can concentrate on their studies (Bettinger et al, 2016). In addition to these financial benefits, there may also be psychological benefits to the receipt of aid. Students receiving funds may treat the interaction as a gift exchange, thus increasing academic effort, or alternatively could gain a confidence boost through the perception that aid has been awarded based on ability (DesJardins et al, 2010).

Understanding the link between financial aid and student outcomes is a particularly important policy question. Socio-economic gaps do not disappear once students enter the door of HE. Low income students are more likely to drop out of college or perform poorly in exams (Bettinger, 2004; Crawford, 2014). And whilst college participation has grown substantially among young people from the poorest background over the past 20 years, completion rates have remained stubbornly low for this group relative to their richer counterparts (Bailey and Dynarski, 2011). The resulting 45 percentage point gap in college completion among rich and poor students, in conjunction with the high wage returns to a degree (Card, 1999), is a concern for income inequality and social mobility.

However, as is widely acknowledged in the literature (Dynarski, 2003; Bettinger, 2004) estimating the impact of student aid on educational attainment is an empirical challenge. There are three main issues at play. First, student aid tends to be correlated with many observable and unobservable factors that also affect an individuals' educational attainment. Aid recipients are more likely to be from poor backgrounds and so are more likely to drop out for reasons unrelated to receipt of aid. Hence, estimates of the impact of aid on attainment are likely to be downward biased. To overcome this issue, researchers would ideally study a form of aid in which students from the same income group receive varying amounts.

A second problem is that aid is often related to student ability. This can arise directly, through merit based programmes, or since more able students congregate at prestigious universities, who may in turn provide more generous financial packages. Since students attending these colleges generally have better university outcomes, this will exert a positive bias on coefficients measuring the impact of aid on student attainment. To overcome this issue, researchers would ideally study a form of aid in which students of the same ability level receive varying amounts.

Finally, the prospect of aid receipt at university may be correlated with the extensive margin in terms of a) a students' likeliness to attend college in general, and b) a students' likeliness to attend a particular college, making it difficult to separate the effects of aid receipt from enrolment effects. To isolate the effect on the intensive margin, researchers would ideally study a form of aid which could not have any influence on a student's decision to attend college, but could only impact their behaviour once they had enrolled.

In this paper we study a particular form of student aid – the English higher education bursary scheme – which we argue has unique features which help to overcome these three problems, and establish the causal impact of this element of financial aid on student performance at university.

As is the case in most student aid systems around the world, poor students are more likely to receive financial aid under the English bursary scheme. However, the definition of a poor student, and how much they receive in aid, varies across institution. In other words, whilst typically there would be no observable counterfactual to a poor student receiving aid (since all poor students receive aid), our data contain a range of counterfactuals at different levels of parental income. The scheme also varies within institution over time, meaning that students from different cohorts within institution receive substantially different amounts of aid. This setup effectively solves the first two of our problems; students with the same parental income and ability can have access to very different amounts of aid.

As well as this variation in aid within university over time, in our set up, aid eligibility is based on strict formulas, which are highly non-linear functions of parental income. Thus, in our most

restrictive specification, we can also exploit sharp changes in aid awarded for a small change in parental income within a university cohort to identify their effects.

The English higher education bursary system is also unique in that it is unadvertised to students, highly opaque, and does not require students to apply to receive it. This, in conjunction with the intricacies and timings of the English university application system, make it near impossible for students to know how much bursary aid they will receive in advance of attending university, making it irrelevant to their decision making (Corver 2010; Calendar and Wilkinson, 2013). This system ensures that bursary aid does not impact both forms of extensive margin, neither the decision to apply to college at all, nor the decision to apply to a particular college. Thus, we can be confident that our estimates of the impact of financial aid relate purely to the intensive margin.

All undergraduate students attending English universities are eligible for the aid scheme, regardless of their age. Moreover, the existence of multiple discontinuities in aid across a wide range of parental income groups means our results are more generalizable than studies which exploit single discontinuities. Thus, our study advances on work by Denning (2016), whose investigation of the impact of aid on enrolled students relates only to older students, and to students with parental incomes around a specific discontinuity. Unlike our work, his estimates also rely on students applying for the scheme, meaning they are more likely to be a highly motivated group and the aid is conditional on minimal standards. Similarly, we also advance on work by Bettinger (2004) who also uses discontinuities in the Pell Grant formula arising from family size to identify the impacts of aid. Again, students must apply for Pell, and moreover, whilst discontinuities in the Pell grant formula are likely to be independent of college choice, they are unlikely to be independent of the decision to attend college.

To examine the impact of bursary aid on college completion and degree performance we collected administrative data from nine higher education institutions in England which contains detailed information on student's finances and attainment throughout college, as well as detailed measures of their human capital upon arrival at university. Our preferred within university-cohort

estimates show that each £1,000 of financial aid that students are eligible for in the first year increases the chances of obtaining a good degree by 3.7 percentage points. This is driven by improvements in both degree completion and course scores. We find that at the mean each £1,000 of financial aid awarded in the first year increases completion by 1.5-1.9 percentage points, and increases test scores by 0.03 to 0.06 standard deviations, depending on the year of study.

The remainder of this paper proceeds as follows. Section 2 outlines the UK student aid system, and the unique features of the higher education bursary scheme. Section 3 describes our dataset. Section 4 outlines our empirical strategy, whilst results, robustness checks and heterogeneity can be found in Section 5. Section 6 concludes.

2. Institutional setup

The UK higher education system is characterised by high tuition fees and high financial support. During the period covered in this paper (student enrolment between 2006-2011), tuition fees for domestic students were capped at £3,000 per year and were binding for all bar a couple of institutions.² These fees were typically not payed up front. Instead students could take an interest free loan³ from the government which covered the entirety of the fee and was paid directly to the university. Students could also take out a maintenance loan at the same rate, with the amount being limited by parental income, location of study (London v elsewhere), and whether or not the student chose to live at home. The highest amount for the lowest income student living away from home,

² The only English universities not to charge the maximum were: University of Northampton, Thames Valley University, The College of St Mark & St John, Leeds Trinity & All Saints (Reddin, 2010). In 2012 fees increased to £9,000 per year, making them the highest in Europe and on a par with many US institutions. While tuition fees are decided at the institution level, in practise the vast majority of universities charge the full £9,000 per year, and the average fee stood at £8,830 in 2015 (See OFFA (2015): Table 2)

³ Although loans were subject to RPI inflation. Real interest rates were added in 2012.

outside of London was £3,555 per year (2006/07).⁴ Both of these loans were repayable after graduation once the graduate was in employment and earning £15,000 per year or more. In addition to student loans the government also provided a non-repayable means-tested maintenance grant of up to £2,700 per year for students with parental incomes less than £17,500 (2006/7).⁵ This amounted to a large level of state support; students starting in 2006/7 with zero reported parental income would receive £6,255 in maintenance loans and grants per year in addition to the loan covering their entire tuition fee liability.

The financial aid program that is the focus of this paper is the English higher education bursary scheme. This was introduced by the UK government in 2006, alongside the raising of maximum tuition fees from £1,200 to £3,000 per year, as a way to placate opponents of the fee increase (Callendar & Wilkinson, 2013). As part of these tuition fee reforms English institutions were required to offer financial aid in the form of a bursary to all low parental income students (defined as those in receipt of a full maintenance grant). The terms and criteria for the minimum bursary – to be set on the basis of parental income, rather than merit or any other criteria – were set centrally by the government and fixed for all institutions. The minimum bursary was set to be the difference between full fee charged and the maintenance grant received by the student. Thus, the minimum bursary for the most deprived students at the time was £300 per year (£3,000 in fees minus £2,700 grants). From 2010 onwards, the rules were redefined so that the minimum bursary became 10% of the tuition fee limit.⁶

⁴ For full-time undergraduates living away from home and studying at English universities outside London. Different rates apply for those living at home or studying in London. See <https://www.gov.uk/student-finance/loans-and-grants> for full details.

⁵ However, maintenance grants will be abolished from September 2016 (with maintenance loans increased to make up the difference), meaning higher education bursaries – the form of aid studied in this paper – will become the sole non-repayable form of student financial aid. Bursaries also represent the sole form of aid which is governed at an institutional rather than a national level, giving rise to significant variation in eligibility across institutions, unlike the other national forms of aid. See <http://www.ifs.org.uk/publications/7905> for more details

⁶ Since 2012, no minimum bursary requirement has been in place. The bursary system was supplanted by the National Scholarship Programme (NSP) in which universities were allocated a set amount of money to distribute among their disadvantaged students in the form of bursaries, fee waivers or other benefits. The NSP has since been disbanded.

In practice, universities typically spent more than the minimum required on these forms of aid, and extended it to more students. Thus, the bursary system became a substantial portion of aid, consisting of £300m in annual expenditure. Around 44% of students receive a bursary, with the average the amount received around £800 per bursary holder per year.

Other than the regulations described above, universities were provided with no guidance or research on how to allocate these funds. Rather, they were given complete independence in how much to give out and to whom. This resulted in universities designing their schemes on a somewhat arbitrary basis, with substantial variation in bursary generosity across institutions, and amounts per holder ranging from as little as £50 to as much as £3,200 per year. Moreover, there is a large degree of cross-cohort variation within institution over time as universities experimented with their schemes from year to year.

This set up effectively solves the first of the identification issues; that aid recipients are more likely to be from poor backgrounds. With the English bursary scheme, the definition of a poor student, and how much they receive in bursary, varies across institutions and within institutions over time. Therefore, while typically there would be no observable counterfactual to a poor student receiving aid (since all poor students receive aid), our data contain a range of counterfactuals at different levels of parental income. This is best illustrated by Figure 1, which shows average bursary eligibility over percentiles of the parental income distribution of students that received bursary aid. As is evident, for students of similar income backgrounds, there is a substantial range of bursaries on offer. For example, students with parental incomes in the bottom 10 percent of the income distribution could receive as little as £350 and as much as £3,200 per year, with variation arising from the university attended and year of entry. Moreover, it is also evident that students across a wide range of the parental income distribution are eligible for a bursary, in one university students from families with income up to £60k per annum, 95th percentile (HMRC, 2016) are eligible for bursary aid.

The second concern is that aid can be related to student ability. Typically financial aid schemes award students with high entry test scores. Hence, we might be concerned that students with

high prior attainment also receive high amounts of aid. In this case estimates of the effect of aid on student performance would be biased upwards. However, as explained, the English scheme is non-merit based. So, conditional on parental income, bursary receipt is orthogonal to entry test scores. Moreover, as Figure 2 shows, whilst universities vary in quality, there is substantial overlap in entry test score requirements. This is driven by universities having different entry requirements for different subjects, particularly arising when a university has a spectrum of reputations across subject areas. Hence, there is a high incidence of common support in test scores across institutions.

However, we may also be concerned that high ability students attend certain institutions which may also offer higher bursaries. Our set up also alleviates this issue by exploiting changes in the bursary schemes within universities over time. By way of example, Figure 3 shows how the bursary schemes of our 9 universities changed between each year of observation⁷, with each panel representing a different university. Looking at University 1 in Figure 3, for example, the maximum bursary that could be received was set to £3,000 in 2006 and then subsequently decreased to £1,000 in 2010, while the maximum parental income of eligible students increased from £15,000 to £25,000 over the same period. Moreover the number of different levels of bursaries awarded and at this university decreased from three to two.⁸ Thereby including university fixed effects along with controls for subject studied and enrolment test scores, we are ostensibly accounting for student ability and exploiting variation within a university across cohorts.

In our most preferred specification, we exploit the sharp changes in bursaries awarded for a small change in parental income within a university-cohort. Again see Figure 3, which highlights the discontinuities in bursary aid awarded within universities according to parental income. Here, we

⁷ Note, whilst our period covers university entry between 2006-2011, some of our universities are left or right truncated

⁸ In practise, for around half of institutions in our sample, students are subject to the bursary rules in place upon year of entry to the course, so that policy change occurring during the duration of their course do not affect them, only new entry students. However, for the remainder policy rule changes affect all students regardless of entry year.

identify the impact of aid through imposing a smooth relationship of outcomes with parental income coinciding with sharp changes in aid eligibility.

Finally, our estimates of the intensive margin could still be biased if students choose to go to university because of the financial aid on offer, or indeed, select their institution based on its financial aid package, knowing that they will be eligible for a certain amount of aid – the third of our identification problems. For example, students from poor backgrounds might choose a particular institution if doing so would mean they gain from a particularly generous aid package, and may also be more likely to graduate for unobservable reasons unrelated to the generosity of aid e.g. motivation. However, in our setting, students are unlikely to sort on parental income, conditional on entry test scores, because they have very little possibility of knowing what their bursary is likely to be ahead of enrolling in college. This non-salience effectively means that the aid program cannot a) affect students' decision to go to university, or b) affect their choice of university. Thus we can be confident we are estimating the impact of aid on the extensive margin only. This non-salience arises for three main reasons, which we now explain in depth, given their importance to our findings.

First, each university has its own unique bursary scheme in place which typically changes on a yearly basis. Universities are required to submit their bursary scheme to the Office for Fair Access (OFFA), in order to show they are complying with government regulations, but are not required to advertise the scheme⁹, meaning schemes are rarely advertised in university prospectuses or included in aggregated university guides. This may not be surprising as universities have an implicit disincentive to attract poor students to attend, since this will then result in increased aid expenditure, as well as having potentially negative peer effects. Thus, in order to know what they are entitled to, students would have to actively search through the finance pages of each institution of

⁹ Universities must report their bursary offer to the Office for Fair Access (OFFA), the English fair access watchdog, in order to satisfy requirements that they are making efforts to widen participation. Universities that do not satisfy OFFA that they are taking steps to attract poor students may not charge fees above £6,000 per year. In practice, no university has ever been refused the right to charge higher fees.

interest, or alternatively, be aware of the OFFA website, download the OFFA agreements and search the documentation for information relating to bursaries.¹⁰

Second, students face a large uncertainty at the time of application about which university they will eventually attend, and hence, which bursary scheme would be applicable to them. In the autumn before students intend to attend university, they apply to up to six universities through the Universities and Colleges Admissions Service (UCAS) a centralised application system, by submitting their predicted A-level entry scores and a personal statement (alongside payment of a fixed fee). In the spring students receive conditional offers from each of their chosen universities. At this point students must rank their top two conditional offers (most favoured and a backup university), with this decision recorded by UCAS. Students then take their A-level examinations in May, receiving their grades during August, at which point they will know which of their ranked conditional choices they have been accepted to. The institutions are committed to accept the applicant and the applicant is committed to going, if they meet the A-level requirements. Only once the student has been accepted at an institution do they then receive a letter informing them of the bursary aid they will receive.

The upshot of this process is that students could not know ahead of accepting their university place what bursary they will end up with, since they do not know which university will be attending until late in the process. Moreover, not knowing which bursary scheme will be applicable to the student will mean that they cannot game the system by mis-reporting parental income.

The third factor contributing to the non-salience of the bursary scheme is that participation in the scheme is passive, that is, students do not have to apply for a bursary in order to receive one. When applying to university courses in the autumn, in addition to including their predicted entry test scores, students can also include household parental income if they wish to apply for any of the

¹⁰ To remedy this situation, Murphy and Wyness have recently collated the complete set of financial bursary rules for English universities and hosted a simplified version on the Guardian newspaper website for prospective students' use. This is available at <http://www.theguardian.com/education/2015/jun/10/which-universities-offer-the-best-bursaries>

government funded student finances (fee and maintenance loans and grants). Around 90% of undergraduate students choose to do so (Bolton, 2016). This income information is not passed on the universities, but sent to the Student Loans Company (SLC), whose responsibility is to administer student finances.¹¹ The SLC also administers the bursary program; each university supplies SLC with their bursary schedule, and in combination with the parental income information, SLC calculates how much each student will receive, making the bursary payment directly into the students bank account along with any maintenance loan or grant that is due to them, at the start of the first term of university.¹² This means that it is possible that students are not aware that they are in receipt of a bursary as any payment received would be a combination of maintenance grant, loan and bursary. Not only does this substantially reduce the salience of bursaries from the students' point of view, but it also means our estimates will represent a treatment effect as students are defaulted into the system, rather than an intention to treat which is more common in the literature.¹³ Finally, as the SLC validates the parental income with the government tax department, this again reduces the likelihood of gaming by students.

There is also empirical evidence that bursaries do not impact the enrolment choices of students in England. Calendar and Wilkinson (2013) survey students who enter English universities in 2008, coinciding with the mid-point year of our sample period. In accordance with the above, they point out that “[students] are notified [about bursaries] only after they accept a place, when it is too late to inform their entry decision and HEI choice”, and “A third of students surveyed had not yet been told whether or not they would receive a bursary, despite the fact that they were surveyed in October 2008 and had started their HEI course, or were about to.”

¹¹ Universities can opt to administer their own scheme but the vast majority choose to do so through the SLC

¹² Similarly, the offer and bursary letters received by the student appears to be from the university the student is attending, but actually comes from the SLC.

¹³ A further important implication is that take-up of bursary each year is not endogenously related to eligibility in the previous year. For example, it may be the case that students who receive large bursaries in first year (and who may also be more likely to be low income) may be more likely to take up their bursaries the next year (and vice versa). As take-up is not governed by the students' wishes, our results do not suffer from this bias.

More substantive evidence that English bursaries do not influence enrolment choices of students has been provided by Corver (2010) who looks at the impact of bursaries on application to university using detailed administrative data. Corver uses individual-level UCAS data on students' choice of university at the time period when students rank preferences of conditional offers. He finds no influence of bursary eligibility on student choice.

Given the empirical evidence and institutional details, we are confident that bursaries do not, and indeed cannot, impact on the extensive margin. But bursary aid may still influence the intensive margin, by impacting students' likeliness to complete each year of their degree, achieve higher university course scores and achieve a good quality degree. This is where we now turn our attention.

3. Data and institutional compliance

This paper makes use of a unique administrative dataset collected from 9 UK universities. The data comprise the entire undergraduate population of UK and EU students for up to 6 cohorts of students beginning their studies between 2006 and 2011. In order to obtain this data we contacted all 118 higher education institutions in the UK, asking them for individual level student data on attainment, parental income and bursaries awarded. Of these 50 agreed to share their data, and we finally received data from 25 all based in England, giving us a sample of 341,398 students. As our estimation strategy relies on using the variation in financial aid for a given level of entry test scores to estimate the effect of bursaries on student outcomes, we first discarded universities who did not provide parental income or those at universities that only provided banded parental income. This reduced our sample to nine universities. We then discarded those students undertaking vocational courses or those above or below degree level. This reduced our sample substantially, leaving us with 35,879 students.

Our sample is truncated, meaning we observe students who started in the earlier cohorts all the way through their studies (3 years), whilst we can only observe the first or second year of students who started later, since they would not have had the chance to complete their degrees at the time we obtained the data. Thus, in our preferred specifications, we use only the non-truncated sample of

students, for whom we are able to observe their full transition through college, including dropouts. This is a total of 22,770 students. In a robustness check, we estimate the impact on outcomes from the first two years of university for students who have enrolled in the appropriate amount years, including continuing students i.e. the full sample of 35,879 students for completing the first year.

The dataset tracks students throughout the course of their degree. Therefore we have information on each student's final degree outcome, including whether they dropped out, their year of drop out, and their average annual course scores. Since these scores are not comparable across universities or individuals, we standardised test scores by university subject and year.

In Table 1 we present some descriptive statistics on the individuals in our full and main samples. The average bursary during the first year per student is £775. To compare bursary aid to state aid that is intended to relieve liquidity constraints, in 2006/7 a student with zero reported parental income would receive £6,255 in maintenance loans (£3,555) and grants (£2,700). In our sample the average bursary aid this student would also receive is £1,138. Average parental income in our sample is £23,288 – though in our sample, we only observe the parental income data of students who provide their data to the SLC (see Section 2, p14) for means-testing of student maintenance loans and grants. Since the upper limit for means-testing is £50,000 we generally observe the parental income of those at or below this limit. This does not bias our estimates since our empirical strategy relies on comparing bursary amounts within bursary holders, the vast majority of whom have parental incomes below this amount.

In regards to the characteristics of the students themselves, 43% of the main sample are male, the average age is around 20 (75% are under 22) and 78% of the sample are white. According to the most recent Higher Education Statistics Agency (HESA) statistics (HESA, 2015: Table 6a), 45% of full-time undergraduates are male, 61% are under 22, and 60% are white. Therefore, our sample is representative in terms of gender, but is younger and whiter than the UK undergraduate population, which is expected given our sample restrictions. Students receive on average £2,035 per year in non-

repayable grant aid from the government. This implies the average bursary received among our sample of students is substantial, representing an additional 38% on grant aid.

In our sample the university completion rate is 90% meaning that only 10% of students fail to complete university. This is compatible with the dropout rate from UK official statistics, of around 8% (HEFCE, 2013) bearing in mind our stricter (degree students) and poorer (income under £50,000) sample. Drop out is highest in first year, at 95%, and steadily declines.

Our main outcome measure is whether a student obtains a *good degree*. We define a good degree as a student graduating from the course with at least an upper second class degree. Unlike the US, students in England rarely drop out from college, however this means many students graduate with low marks. To differentiate students in the subsequent labour market much emphasis is placed on the final grade of the student's degree. The possible grades awarded are Fail, Third Class, Lower Second Class, Upper Second Class and First Class degrees. We define students obtaining a good degree as those being awarded a First or Upper Second Class degree, which equates to 63 percent of all first year enrollees. As a point of comparison the six-year graduation rate for students who started in the fall of 2006 was 60.5 percent at public four-year colleges, and 62.5 percent at private nonprofit colleges (Shapiro et al, 2014). English Graduates with a First or Upper Second Class degree experience a significant wage premium compared to other graduates (Feng & Graetz, 2013; Walker & Zhu, 2013).

Despite the strict institutional setup described in Section 2, we observe a degree of non-compliance in our data. This is illustrated in Figure 4, which plots household income and bursary eligibility versus receipt for our 9 universities in one particular year (2008). As can be seen in this figure, the vast majority of students receive the bursary amount that corresponds with their observed household income. However in a small but significant number of cases, students receive more or less than they are entitled to. Across all our universities, we observe varying rates of non-compliance, with the average of around 5% of students receiving a bursary that is “too high” and around 7% receiving a bursary that is “too low”.

One concern is that these issues are not simply random measurement error, but are arising from systematic issues that could generate biases. Administrators at these universities stated three situations where the amount of bursary received does not equal that which should be received for that level of parental income. First, a reassessment of parental income indicated that the student would be eligible for a different student aid amount (either due to student error, or a sudden change in circumstances). Assuming that the measurement error may have overstated or understated parental income, this would downward bias the estimates. The second type of non-compliance concerns student pre-dropout. If students register for a course, but then withdraw from the course before arrival, they will not receive a bursary but may still be recorded in the administrative records. Typically, such students would have been removed from the data, but it is possible that they could still appear as receiving zero bursary and dropping out in year one, which would bias our estimates upward. The third example of non-compliance concerns the university using its discretion to award additional funds to some students. If it is the case that institutions are systematically awarding high ability students more than they are entitled to, this will again bias our estimates upwards.¹⁴ In order to eliminate the biases caused by this non-compliance, we adopt an instrumental variables (IV) methodology for our preferred specification using the amount of aid that the student is eligible for as an instrument for the amount of aid received. This is described in more detail in Section 4 below.

4. Estimation strategy

Our empirical strategy exploits two sources of variation in financial aid awarded to estimate the causal impact of aid on outcomes. First, we exploit variation of bursary aid for a given parental income within university over time, which arises due to the changes to bursary schemes within institutions (best illustrated in Figure 3). We effectively compare two individuals of the same parental income

¹⁴ Despite students' prior test scores being uncorrelated with indicators of whether the students more or less than their designated amount, one may still be concerned that those receiving more may have other unobservable positive abilities.

background, but receiving different bursary awards due to their university entry year. Given the opaque nature of the student aid system for students applying to university it is unlikely that students delay or bring forward their enrolment at a specific institution in response to changes in bursary schemes.

A second source of variation is more restrictive. We purely exploit the non-linear nature of the bursary schemes within an institution entry cohort, which is again seen in Figure 3. As can be seen, there are sharp discontinuities in the amount of aid awarded for only small changes in parental income. For example, in University 1, we can see an individual with parental income of £15,000 in 2006 would have received a bursary of £3,000, but an individual with parental income of only £1 more would receive a bursary of £1,545. Identification comes from exploiting this highly non-linear relationship between aid received and parental income.

In this specification we include university-cohort fixed effects, whilst also accounting for up to a third order polynomial in parental income, and nonlinear controls for entry test scores, student characteristics and subject area studied. This method ensures that student outcomes should vary smoothly with parental income and so we attribute any remaining non-linearities to the causal effect of bursaries.

Our empirical strategy is somewhat similar to the “heavily parameterized regression kink design” (Clark & Del Bono, 2016), which in turn is analogous to “regression kink design” (Card et al. 2012, Dong, 2010), but requires a stronger assumption, specifically that our third-order polynomial in income must capture the underlying outcome-parental income relationship across the full range of scores. Following Clark & Del Bono, we provide two validity tests of this assumption. First, we test if there is an “effect” of bursaries on pre-determined characteristics, in this case, university entry scores. This tests to ensure that there are no other sharp discontinuities that may be generating the effect (e.g. ability). Second, we check our estimates are robust to alternative polynomial specifications.

In addition to these assumptions, we need to satisfy those required for standard regression discontinuity designs; that students are similar each side of the cutoffs and students cannot dictate the treatment status by ‘gaming’ parental income (Lee & Lemieux, 2010). Given the institutional setting there are three strong reasons to believe that students are not sorting around these cutoffs (as described in more detail in Section 2). First, students don’t know which university they will be attending until very late in the process. Two, students are not aware of the bursary rules due to the opaqueness of the schemes, and the fact that they need not even apply for bursary aid in order to receive it. And third, the parental income submitted to the SLC is validated by the government, therefore making it hard to cheat without consequence. Regardless, in addition for checking for discontinuities in predetermined characteristics, we also check differential densities each side of the cutoffs using Cattaneo, Jansson, and Ma (2016) density tests.

Whilst our estimates, in line with standard regression discontinuity designs (Jacob and Zhu, 2012) will generate “LATE” estimates, i.e. the impact of aid for students close to the discontinuities, the presence of multiple discontinuities across much of the parental income range allows us to use all of the data to help identify the effects of interest, and results in a more generalizable effect than that generated by estimations around a single discontinuity.

As highlighted in Section 3, some universities have not complied with their own stated bursary rules, and so we are concerned that our results may suffer from biases caused by non-compliance. To account for this, we adopt an instrumental variables approach. The estimation consists of two-stage least squares instrumenting the actual aid amount awarded with the amount of aid the student is eligible for, according to the university’s own rules. In the first stage, the size of the aid coefficient therefore represents the average increase in aid the student is eligible for, rather than that awarded. The second stage estimates the relationship between students’ aid eligibility and the outcome of interest. Specifically we use the following equations:

$$y_{ijt} = \beta_1 \widehat{Aid}_{ijt} + f(Inc_{ijt}) + \beta X_{ijt} + \delta_{jt} + \varepsilon_{ijt} \quad (1)$$

$$Aid_{ijt} = \mu(Eligible) + f(Inc_{ijt}) + \beta X_{ijt} + \delta_{jt} + \varepsilon_{ijt} \quad (2)$$

where y is the outcome of student i attending university j , who started in year of entry cohort t . Aid is a continuous variable representing the amount of financial aid received by student i in thousands of pounds. Parental income Inc is accounted for with a third-order polynomial. The detailed nature of the data also allow us to condition on a large vector of background characteristics of all undergraduates in the study (X), such as university entry grades, age, ethnicity, gender and subject of study. In keeping with our estimation strategy, we control for these characteristics in the most flexible way possible, using dummies for each age, gender, ethnicity type, and for university entry grades (the latter variable spans from 0-300, therefore we have a series of 30 dummies for each 10 point range in entry grades). We additionally control for the national student financial aid award (maintenance grants), which is means tested but differs from bursaries since it is awarded at the national rather than institutional level, hence has no across university variation, and only has two kinks.¹⁵ Finally in our most demanding specification we include a set of university-year effects (δ_{jt}), which will provide us with the parameter of interest β_1 the impact of an additional £1,000 of financial aid on student outcome y , exploiting the nonlinear jumps in bursary awarded for a small change in income within an institution year group. We also present estimates that use the within university variation over time, by including the institution and the year dummies separately.¹⁶

5. Empirical results

At this point in a standard RD paper, one would traditionally plot the relationship between the outcome of interest and the running variable, on either side of the discontinuity. However, in our case, not only do we have multiple discontinuities, but in 4 of our 9 universities, the aid schedules have slopes. We instead present the residuals from our main specification (equation 1) *excluding* the

¹⁵ Excluding the national grant scheme in the set of student characteristics does not significantly alter any of the results.

¹⁶ All standard errors presented are robust and clustered at the university-year level

bursary parameters, which we plot against distance from nearest cut-off, only for cut-offs which would generate discrete changes in bursary amounts over £400. This comprises 15 cut-offs over all cohorts and universities, which we stack into one discontinuity. As in our main specification we have a flexible functional form for parental income, and all other characteristics have been dummied out, so we would expect a smooth relationship between income and residuals; any jumps in the residuals around a cut-off is indicative that bursary aid has an impact. Figure 5 plots the residuals from households with income up to £10,000 away from their nearest cut-off. Here we see that students below the cut-off have higher residuals than those just above the cut-off, providing suggestive evidence of a positive impact of bursaries, albeit with a much reduced sample size ($n=3,544$).

We use a manipulation test proposed by Cattaneo, Jansson, and Ma (2016) (CJM), building on the McCrary (2008) test, to test for students gaming the system. This would be seen by there being a higher concentration of students immediately below these discrete cutoffs. The CJM test is a data-driven approach based on a local polynomial density estimator, which does not require pre-binning of the data and automatically specifies bandwidths each side of the cut-off. For this sample of individuals, the CJM test proposes bandwidths of £1859 and £1850 to the left and right of the cut-off respectively. This leaves an effective sample of $N^-=424$ and $N^+=400$ for treatment and control groups respectively. The manipulation test value is $T=0.827$, with a p-value of 0.408, and therefore as expected due to the nature of the institutions, there is no statistical evidence of systematic manipulation of the running variable or sorting across universities.

We also present the same residuals, but this time, using our entire sample, rather than just for individuals near to large discontinuities, plotting them against bursary aid, in Figure 6. This non-parametric means of displaying the results shows that over the first £500 of bursary aid, students perform better than those with zero aid. For bursaries between £500 and £1500, however, the effect of bursary aid levels off. For bursary levels above £1,500 the confidence intervals become extremely wide, since few individuals receive bursaries at such high levels. Given this observed relationship, we allow for decreasing marginal returns to aid in our specifications.

We now use the empirical strategy described above to estimate the causal effects of means tested aid on enrolled students' outcomes such as obtaining a good degree, completing the three years of university and test scores in each year.

A. Obtaining a 'Good Degree'

Table 2 reports our estimates of the impact of bursary aid on the probability of obtaining a 'good degree'. Column 1 presents simple correlations between aid and outcomes, with each successive column introducing additional parameters; parental income, student characteristics, university effects, year effects, and ultimately university-year effects.

In Panel A, we assume constant returns to financial aid, whilst in Panel B we allow for decreasing marginal returns by adding a quadratic term in aid where we report the effects for aid and aid squared. For ease of interpretation, we present the marginal impact of £1000 of aid at the mean level of year 1 bursary aid (£775). Here, the raw correlation between aid and outcome implies each £1000 of aid at the mean results in a 0.5 percentage point reduction in students' chances of graduating with a good degree (column one, Panel B).

There will be both positive and negative biases at play here. On the one hand, students from low income households are more likely to receive more financial aid, and are also less likely to achieve a good degree, generating a negative bias. On the other hand, students with high ability are likely to perform well at university, and are also more likely to attend prestigious richer institutions, which can afford to give out bigger bursaries; these factors would generate positive biases.

In column two we address the first of these issues by controlling for a third order polynomial in parental income. As expected, this greatly raises the marginal effect at mean to 0.11. Column three additionally accounts for student characteristics (age, gender, ethnicity) including entry test scores, which reduces the marginal impact to 0.077. The next columns introduce university and year fixed effects in turn, to account for any further differences across universities or for any general increase in the probability of achieving a good degree over time. The marginal effect remains stable with the introduction of these parameters meaning that the flexible functional form of parental income and

discrete dummies for prior test scores are accounting for much of these potential biases. By column five we are using the variation in bursary aid schedules over time within universities, and the non-linear relationships between aid and parental income as all other characteristics have been dummied out.

The final column (6) replaces the university and year effects with a set of indicators for each year university combination. This only exploits the non-linear relationship between aid awarded and parental income and is our most restrictive specification. Here we find a £1,000 increase in bursary aid at the mean increases the probability of gaining a good degree by 7.1 percentage points.

As described in the previous section, despite bursary aid being administered by the SLC, we observe some non-compliance in aid awarded. To address for the potential of endogenous non-compliance we repeat the above estimations using the bursary aid eligibility, instead of aid received. These parallel set of results are presented in Panel C, where we find the marginal effect at the mean to have over halved in size, to 3.2 percentage points. In Panel D, we present our preferred estimates, which instrument the aid received with the amount eligible. As expected these 2SLS estimates are larger than those from the reduce form, however they remain lower than the original estimates. As before there are decreasing returns to amount of bursary aid, with the maximum impact that aid could have occurring at £1906 (Col 6 Panel D). Here, the marginal effect coefficient reveals an increase in the possibility of gaining a good degree by 3.7 percentage points for a £1,000 increase in aid (significant at the 1% level).

As the endogenous variables are aid and its square, instead of showing the first stage estimates and the associated F-Statistics, we present the Shea's adjusted partial R-squared. In accordance with having similar reduced form results, the high values indicate that the aid rules are very good predictors of amount of aid received.

Note the results from our most demanding specification, exploiting only the non-linearities in aid within university cohorts are very similar to those found in column 5 (Panel D), in which we also exploit the variation arising from changes in bursary rules across cohorts within university. This is a

less restrictive specification, but has the advantage of using more variation, and as the results indicate, may be sufficient.

B. Degree Completion and Course Scores

What could be driving this increase in the chances of getting a good degree? We explore this in Tables 3-4 by looking at the impact on completion of each academic year and annual course scores. All effects presented in these tables are the marginal impacts at the mean and are obtained from separate regressions. Table 3 first shows the impact of an additional £1,000 bursary award in first year on completion of the first, second and third years, along with good degree for comparison purposes. Note students must complete each year to obtain a good degree. In each case, the full set of controls and a quadratic in bursaries is used. For our preferred 2SLS estimator (Panel B), we find evidence that bursary aid has a positive impact on completion. We find a £1,000 increase in aid improves students' likeliness to complete the first year of the degree by 1.4 percentage points.

The impact of aid on completion increases with each year (1.6 2nd year, 1.9 3rd year). Intuitively it would be difficult for these effects to decrease, as if £1000 of aid increased the probability of completing the first year by 1.4 percentage points, then it will have at least this impact on students ever completing the second year, unless second year aid has a negative impact. Therefore these larger coefficients represent the marginal impact of aid. Note this will be a combination of aid receipt in the first and second year, as these aid amounts are highly correlated.¹⁷ These effects are comparable to those found by Bettinger (2004), who finds that a \$1,000 (£660 approx) increase in Pell aid corresponds to a 4 percentage point reduction in the likelihood that students withdraw from college in first year, given exchange rates and inflation.

The impact of aid on completion of final year is lower than the impact of aid on obtaining a good degree – suggesting there may be some additional impact of aid coming through course scores.

¹⁷ Estimating the additional impact of aid in the second and third years proved to be difficult given the high collinearity between bursary aid received across years, and the possibility of students gaming reported parental income subsequent years. This is despite the fact that in some some universities change the aid schedules as students progress.

This is examined more fully in Table 4, in which we present the results for an additional £1,000 of bursary on mean standardised course scores each year. Here we see a largely positive impact of bursaries – with the IV estimator showing an additional £1,000 of bursaries in the first year generating a 0.064 standard deviation increase in course scores in that year, a 0.041 standard deviation in course scores in the second year, and an insignificant impact in the third year of 0.03

In summary, our analysis shows a positive impact of bursary aid on obtaining a good degree, to the tune of 3.7 percentage points per £1,000, against a mean good degree rate of 62 percent. This positive impact appears to be driven by both an increased probability of completion (of as much as 1.9 percentage points) and improvements in test scores (of as much as 0.064 standard deviations). These impacts are comparable to those obtained by Bettinger (2009).

6.2 Robustness Checks

We perform a series of robustness checks on our IV estimates to determine their stability. These are shown in Table 5. The first row again presents the marginal effects of aid in the first year of study at the mean for our preferred IV models and outcome measures. The outcomes are complete the 1st year (Column 1), standardised 1st year course scores (Column 2), and obtain good degree (Column 3).

Our main specification only uses students who could have potentially completed their course. However, we have data on all students that are currently studying at these 9 universities. Therefore the second panel shows estimates includes additional cohorts, including all current students (i.e. those for whom we can only observe to the end of first or second year), this increases the sample size by around 13,000 to 35,879. Reassuringly the estimates are very similar when including these additional observation. We do not present estimates for Good Degree as the continuing students have yet realise this outcome.

One of the arguments that we put forward is that comparisons can be made across universities, which we support by showing that there is common support in the entry test scores of students (See

Figure 3). Three universities appear to be exceptions to this, university one and eight appear to only enrol students with the highest of test scores, and university two appears to enrol students whose test scores are mostly below that of the others. Therefore in the third row we re-estimate the results excluding these universities. Again this appears to have very little effect on the results.

Similar to a regression discontinuity paper, we want to establish that treated individuals are similar to the untreated. In the case of a heavily parameterised regression kink design, we want to show that bursary receipt is uncorrelated with any pre-determined characteristic. To do this we estimate the impact of bursary aid on entry test scores using our main specification, omitting the entry test score controls. These results are shown in the fourth column, and we find that there is no significant relationship using any of the sub-samples.¹⁸

The other critical assumption needed for this estimation strategy is that the functional form of the running variable is sufficiently flexible, that with the highly parameterised specification the relationship between parental income and the outcome is smooth. To test this we present estimates with alternative polynomial specifications of parental income, from linear up to quintic, on first year test scores, completion and obtaining a good degree. These can be seen in Table 6 for and show that our results are robust to the order of polynomial, and the cubic relationship we use throughout is sufficient.

6.3 Heterogeneity

We now consider whether the relationship between bursary aid and outcomes varies according to student characteristics by re-running our main model (equation 2) for different groups according to gender, ethnicity, parental income and prior test scores. These results are presented in Table 7.

As can be seen, we see little difference in the impact of bursary aid by gender. The marginal effects are similar for completing the first year, course scores, and obtaining a good degree – although, for the first of these the estimates are only significant for males. For age we see there are distinct and

¹⁸ Falsification tests on the other pre-determined characteristics have been run and are available upon request.

significant differences in the impact of bursaries according to the age of the student. The positive impact of bursaries appears to be driven solely by traditional age students (those who enter university at a young age, less than 20 years old), rather than more mature students.

A common concern among higher education policymakers is that that need based aid may simply subsidize college for infra-marginal students – those who may not benefit from college but are induced into it due to the low costs (Dynarski, 2003). While that paper was concerned with the extensive margin, we can directly test which types of students do benefit from our aid package, in terms of both disadvantage, and ability. By doing so we be able to estimate the impacts along two dimensions that a student may be marginal, ability and liquidity. In each case, we run regressions split by income and ability group, defined as above or below the median of all students. For ability we also create a “relatively high ability” and “relatively low ability” students, here defined as above the median of all students within the student’s university of attendance.¹⁹ We believe these findings will be more informative for forming individual university aid policies based on merit.

First, in terms of parental income, we find that the poorer students in the sample gain considerably more than richer students. The estimated impacts for students in the lower half of the distribution are around six times higher than estimates on the whole sample for all outcome measures. This suggests that means-based aid is not simply subsidizing infra-marginal students in terms of ability who would not gain from university experience, but actually acting to improve their outcomes at university. Turning to ability directly, we find that the benefit of aid is much higher for previously high achieving students. Students from the top half of the prior achievement distribution gain two to three times more from bursary aid than those from the bottom half of the distribution, in terms of test scores, completing the first year and obtaining a good degree. This suggests that there are some high ability students facing liquidity constraints, which bursary aid is acting to relieve. Looking within

¹⁹ We do not provide estimates for relatively high/low income by institution as the estimation method relies on deviations in the outcome based on predictions from parental income within institution. Splitting the data within institution by parental income will improve the fit of the line and reduce the variation that can be exploited.

institutions, we find only significant impact of aid on those in the top half of the prior achievement distribution. This has a further implication that university aid packages which target high achieving students may be more effective than those purely based on a means-test.

7 Conclusion

The majority of studies of the effectiveness of student aid focus on its effects on enrolment. This paper instead examines the causal effect of aid on the outcomes of students who have already enrolled in a college. To do so we exploit changes and nonlinearities in financial aid schedules as sources of variation, and find that unconditional financial aid in the form of an annual bursary increases students' likelihood of obtaining a good degree by 3.7 percentage points for each additional £1,000 of aid awarded. This positive effect is driven by both improvements in test scores and in degree completion rates. Given the unique institutional setting of bursary aid – that students are likely not aware of a university's scheme until a month before term starts and they are committed to attend– we are confident that these results relate purely to the intensive margin, rather than being driven through enrolment effects. Therefore the contribution of our paper is to provide rare evidence on the impact of aid on the intensive margin. These effects are comparable to those generated by Bettinger (2004), who finds a 4 percentage point effect on dropout for \$1,000 of Pell Grants.

Our finding of a positive impact of aid on degree performance has several policy implications. Most importantly, our evidence suggests that aid matters even once students are enrolled in college, helping them to obtain better class degrees through improving persistence and course scores. Given that our form of aid does not act as a subsidy to the cost of college, it is likely that the additional financial resource provided by bursary aid is directly relieving liquidity constraints faced by students.

However, universities and policymakers should not simply assume that all aid to all student types will be equally effective. We found that the marginal impact of financial aid is decreasing and from our fully parameterized estimation with a quadratic in aid, we find that aid continues to have a positive impact up to £1900.

There are two issues that policy makers need to bear in mind when designing the optimal aid package – liquidity and ability. First, the ideal aid package should encourage liquidity constrained students into college by lowering the cost of university to such a place where college becomes affordable, and to provide enough financial support to enable them to succeed. Second, ideally aid should not simply subsidize infra-marginal students in terms of ability – those who may not have the necessary skills to succeed at college, but are drawn in due to the lower cost. Equally, it should not simply act as a transfer payment for ultra-marginal students; those who would attend college and do well regardless of the financial aid on offer.

Although we do not examine enrolment effects in this paper, we can consider the impact of aid received at university, bearing in mind these two issues, by looking at students in terms of ability and liquidity.

Since our results indicate that aid receipt does have a positive impact on persistence and degree performance, this is evidence that aid is not simply subsidizing the infra-marginal student (based on ability); our results show that students receiving bursary aid persist further and achieve better outcomes than they would have done without the subsidy.

Our results also show that students from richer backgrounds gain less than those from poorer backgrounds. Therefore an efficient use of aid resources would be to attribute more to the lower income students and less to the high income students (who could be considered ultra-marginal in terms of liquidity). Whilst this is encouraging for proponents of means-tested aid, we should also consider that aid packages which are exclusively means-tested may not be the most efficient use of societal resources. Our results also show that high ability students benefit the most from aid, suggesting an important role for a merit component. This type of aid, that is dependent on merit and demonstrated need, is more common in the US in the form of scholarships, but less so in the UK.

What do these findings tell us about the effectiveness or otherwise of current aid packages at UK institutions. Universities give out, on average, £775 of bursary aid to each student, though with a large degree of variation around this, suggesting that many institutions could streamline their aid

policies to improve efficiency. For example, many of the highly selective English institutions give out large aid packages to a small number of students (e.g. the poorest students at Imperial College, ranked 8th in the world²⁰, receive £6,000 per year in bursary aid). Given the decreasing returns to aid, such institutions should distribute their resources across more students, especially since their students are likely to be highly able, and therefore to gain the most.

Meanwhile, less selective institutions tend to give out smaller amounts of aid to larger numbers of students (e.g. Liverpool John Moores awards bursary aid of at least £400 to 65% of their students). Our evidence suggests they should instead give out more aid to the most able of their students.

This bursary aid policy came about as a result of increasing tuition fees in 2006, to alleviate concerns relating to participation of disadvantaged students, enforcing universities to distribute this additional fee income in the form of bursaries. This parallels the price discrimination that takes place at private non-profit institutions in the US, with the richer students effectively subsidizing the poor through higher fees. We find this aid to have had an impact on the intensive margin. Given the long run labour market impacts of obtaining a good degree (Walker and Zhu, 2013), this suggests this form of cross- subsidization is money well spent.

References

- Bailey, M. and Dynarski, S. (2011) Inequality in Postsecondary Attainment, in Duncan, J. and R. Murnane (eds.) Whither Opportunity, Russell Sage.
- Bettinger, E (2004) 'How Financial Aid Affects Persistence', NBER Working Paper 10242, National Bureau of Economic Research
- Bettinger, E., Gurantz, O., Kawano, L., & B. Sacerdote, (2016), 'The Long Run Impacts of Merit Aid: Evidence from California's Cal Grant', NBER Working Paper 22347, National Bureau of Economic Research
- Blundell, R., Dearden, L., Goodman, A. and Reed, H. (2000), The Returns to Higher Education in Britain: Evidence From a British Cohort. *The Economic Journal*, 110: 82–99.
- Bolton, P (2016) 'Student Loan Statistics' House of Commons Library, Briefing Paper, No1079
- Bound, J. & S. Turner, (2002). Going to War and Going to College: Did WorldWar II and the G.I. Bill Increase Educational Attainment for Returning Veterans? *Journal of Labor Economics*. 20(4): 784-815.
- Callender, C & D.Wilkinson (2013). Student Perceptions of the Impact of Bursaries and Institutional Aid on their Higher Education Choices and the Implications for the National Scholarship Programme in England. *Journal of Social Policy*, 42, pp 281-308
- Card, D. (1999), 'The causal effect of education on earnings', in O.Ashenfelter and D.Card, (eds), *Handbook of Labor Economics*, vol. 3A, North-Holland.
- Card, D., Lee, D., Pei, Z., & A. Weber, (2012), 'Nonlinear Policy Rules and the Identification and Estimation of Causal Effects in a Generalized Regression Kink Design', NBER Working Paper 18565, National Bureau of Economic Research
- Clark and Del Bono, (2016), 'The long-run effects of attending an elite school: Evidence from the United Kingdom', *American Economic Journal: Applied Economics*, 8 (1) (2016), pp. 150–176
- Corver, M., (2010) 'Have bursaries influenced choices between universities?', Research Paper 2010/06, Office for Fair Access, UK

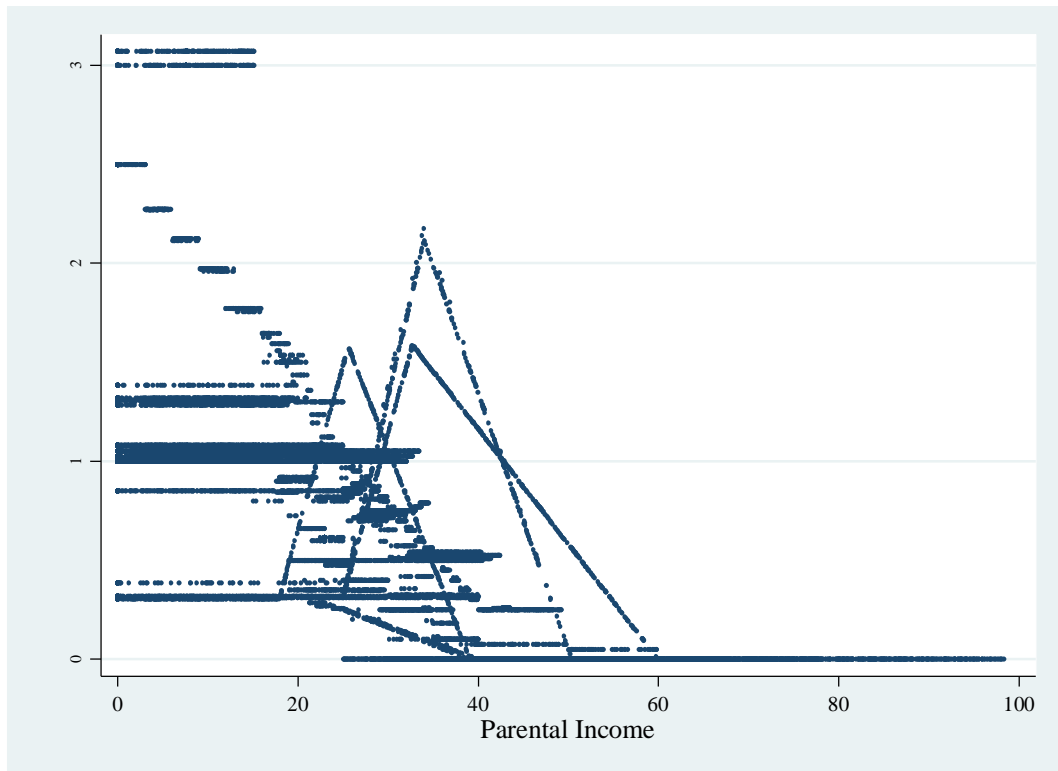
- Crawford, C. (2014). Socio-economic differences in university outcomes in the UK: drop-out, degree completion and degree class. IFS Working Paper W14/31, Institute for Fiscal Studies.
- Dearden, L., E. Fitzsimons, and G. Wyness (2014). ‘Money for nothing: Estimating the impact of student aid on participation in higher education’, *Economics of Education Review*, Volume 43, December 2014, Pages 66–78
- Denning, J., (2016), ‘Born Under a Lucky Star: Financial Aid, Credit Constraints, and College Completion’. Unpublished, Brigham Young University, Provo UT.
- DesJardins, S. L. and B. P. McCall (2010). Simulating the effects of financial aid packages on college student stopout, reenrollment spells, and graduation chances. *The Review of Higher Education* 33 (4), 513–541
- Dynarski, S. (2000) ‘Hope for Whom? Financial Aid for the Middle Class and Its Impact on College Attendance’, *National Tax Journal*, 53, 629–61.
- Dong, Yingying. 2010. “Jumpy or Kinky? Regression Discontinuity without the Discontinuity.” Working Paper No 111207, University of California-Irvine, Department of Economics
- Dynarski, S. (2003) ‘Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion’, *American Economic Review*, 93, 279–88.
- Garibaldi, P., F. Giavazzi, A. Ichino, and E. Rettore (2012). College cost and time to complete a degree: Evidence from tuition discontinuities. *Review of Economics and Statistics* 94 (3), 699–711.
- Feng & Graetz (2013), ‘A Question of Degree: The Effects of Degree Class on Labor Market Outcomes’, CEPDP1221. Centre for Economic Performance, London School of Economics and Political Science, London, UK
- HEFCE (2013), ‘Non-continuation rates at English HEIs’, Higher Education Funding Council for England
- HESA (2015), Students and Qualifiers data tables, table 6a, Higher Education Statistics Agency
- HMRC (2016), ‘Percentile points from 1 to 99 for total income before and after tax’, HM Revenue & Customs, Available at <https://www.gov.uk/government/statistics/percentile-points-from-1-to-99-for-total-income-before-and-after-tax>. Last Accessed 25/07/2016

- Jacob, R. T., P. Zhu, M-A. Somers, and H. S. Bloom. (2012). *A Practical Guide to Regression Discontinuity*. MDRC Working Paper. New York: MDRC.
- Joensen, J. (2013). *Timing and Incentives: Impacts of Student Aid on Academic Achievement*. mimeo, Stockholm School of Economics.
- Kane, T. (1995) 'Rising Public College Tuition and College Entry: How Well Do Public Subsidies Promote Access to College?', National Bureau of Economic Research (NBER) Working Paper 5164.
- Lee, D., & Lemieux, T., (2010). "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, American Economic Association, vol. 48(2), pages 281-355, June
- Leslie, L. L., and Brinkman, P. T. (1988). *The Economic Value of Higher Education*. New York: Macmillan.
- Nielsen, H.S., T. Sorensen, C. Taber, 'Estimating the effect of student aid on college enrollment: Evidence from a government grant policy reform', *American Economic Journal: Economic Policy*, 2 (2) (2010), pp. 185–215
- Organisation for Economic Co-operation and Development (OECD) 2010. *Education at a glance*. Paris: OECD
- Psacharopoulos, G, 1994. 'Returns to Investment in Education: A Global Update,' *World Development*, 1994, 22 (9), 1325-1343
- Scott-Clayton, J. (2011). *On money and motivation: A quasi-experimental analysis of financial incentives for college achievement*. *Journal of Human Resources*, 46, 614–646.
- Seftor, N. & S. Turner, (2002) 'Back to School: Federal Student Aid Policy and Adult College Enrollment', *Journal of Human Resources*, 37, 336–52.
- Sjoquist and Winters (2012), 'State Merit-based Financial Aid Programs and College Attainment', IZA DP No. 6801
- Schudde, L. & Scott-Clayton, J (2014), 'Pell Grants as Performance-Based Aid? An Examination of Satisfactory Academic Progress Requirements in the Nation's Largest Need-Based Aid Program', CAPSEE Working Paper, December 2014

UCAS (2015) 'End of Cycle 2014 Data Resources; Acceptances by provider country and domicile', University College Admissions System, available at:https://www.ucas.com/sites/default/files/eoc_data_resource_2014-dr3_005_01_0.pdf

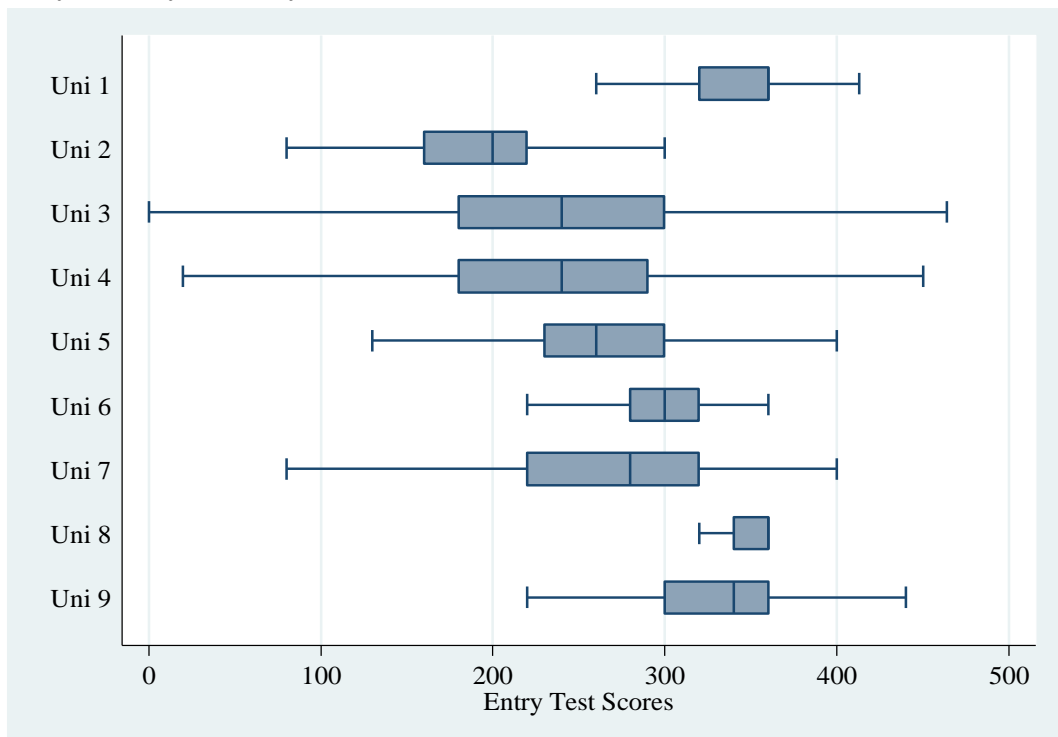
Walker, I. and Zhu, Y. (2013) The impact of university degrees on the lifecycle of earnings: some further analysis. London, UK: Department of Business, Innovation and Skills.

Figure 1: Eligible Bursary by parental income



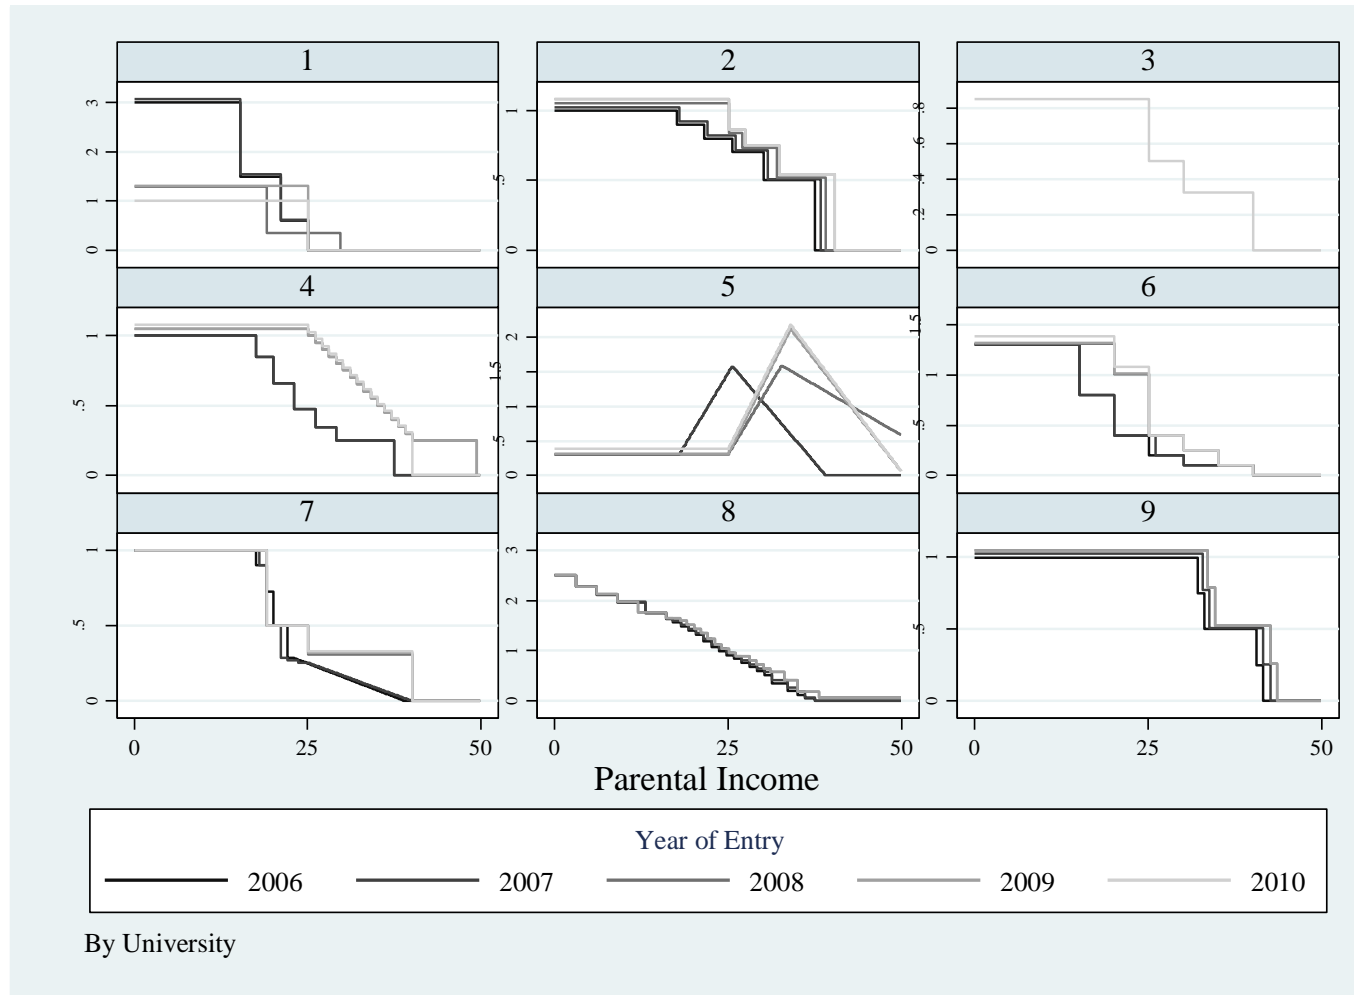
Notes: Each point represents the amount of bursary aid available for first year students at each university and entry cohort. Figures reported in nominal values. Source administrative data from the 9 universities.

Figure 2: Entry scores by university



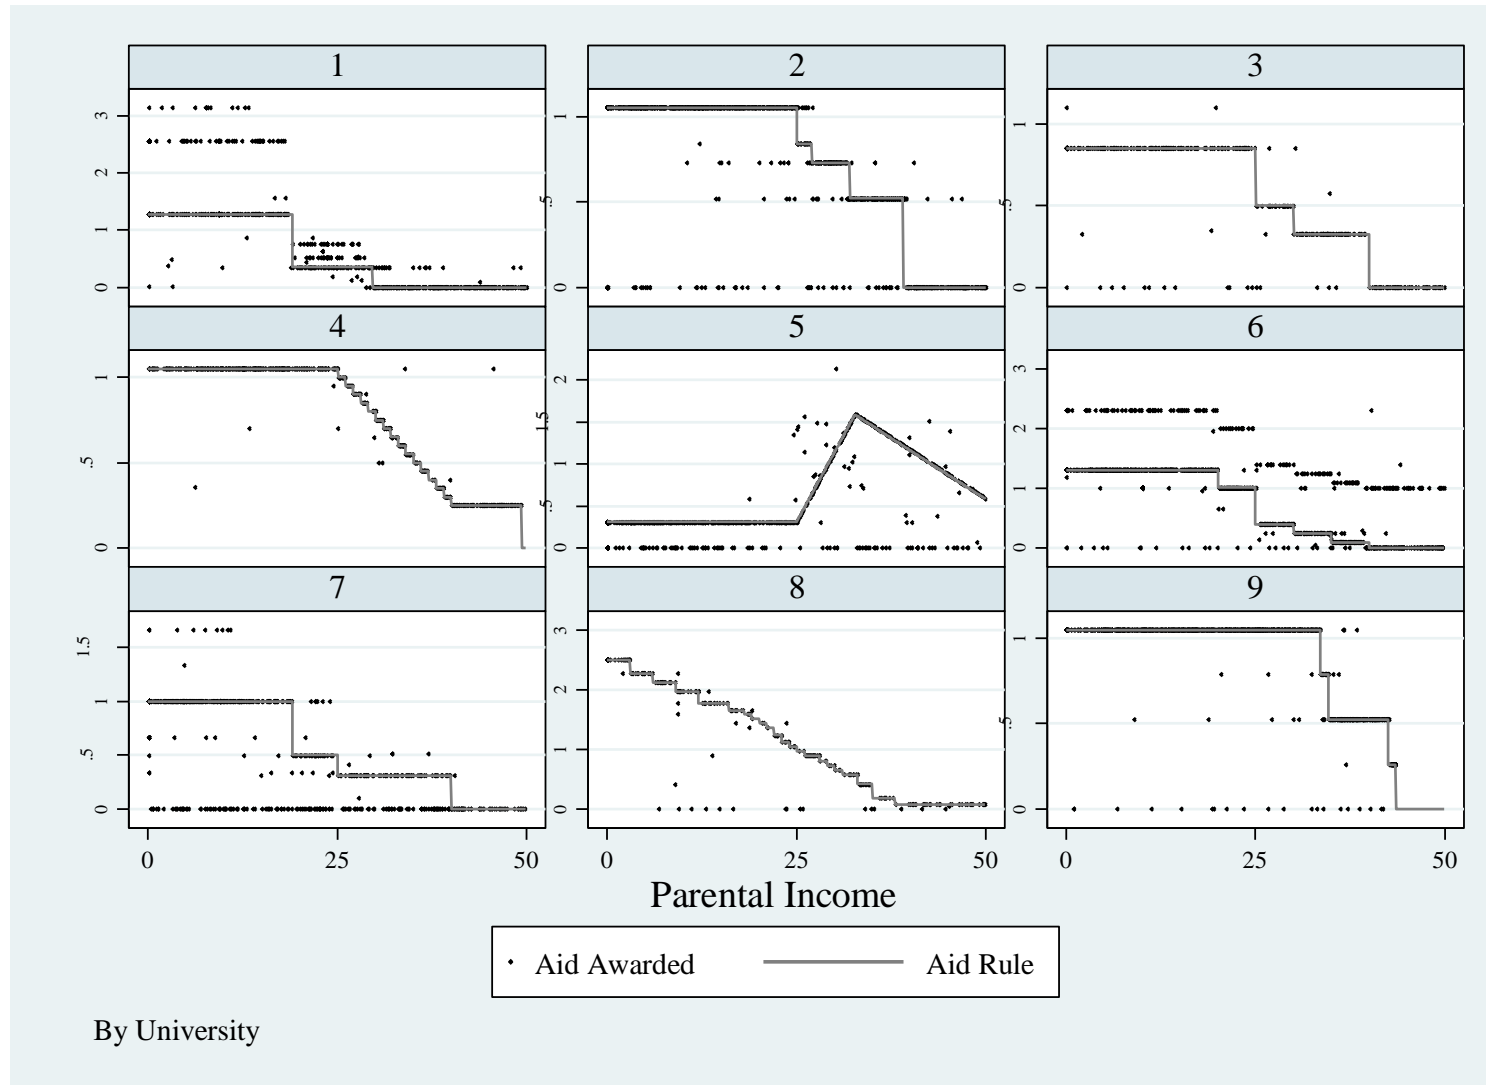
Notes: Figure 2 shows box plots of the entry qualification scores of students attending each university in the estimation sample. The ends of each box represent the 25th and 75th percentiles in entry qualification scores. Source administrative data from the 9 universities.

Figure 3: Financial Aid Schedules at Universities over time



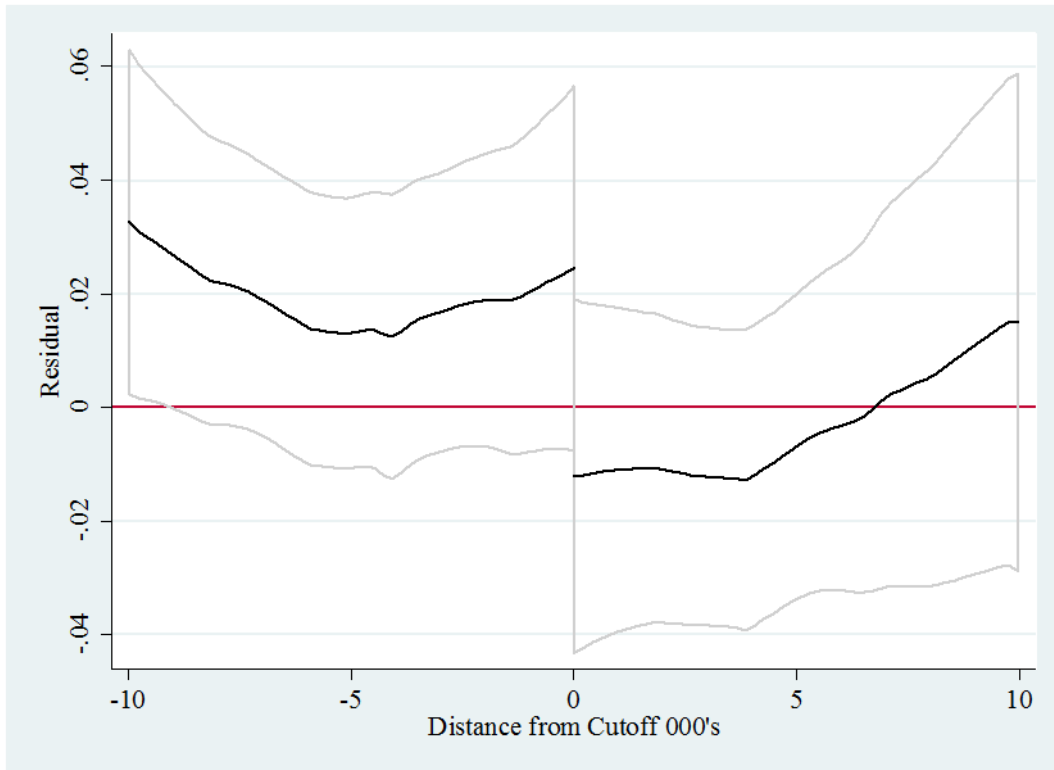
Notes: Represents the financial aid schedules for first year students for nine anonymous English universities for students entering in the years 2006 through to 2010. Figures reported in nominal values

Figure 4: University aid rules and compliance at universities in 2008



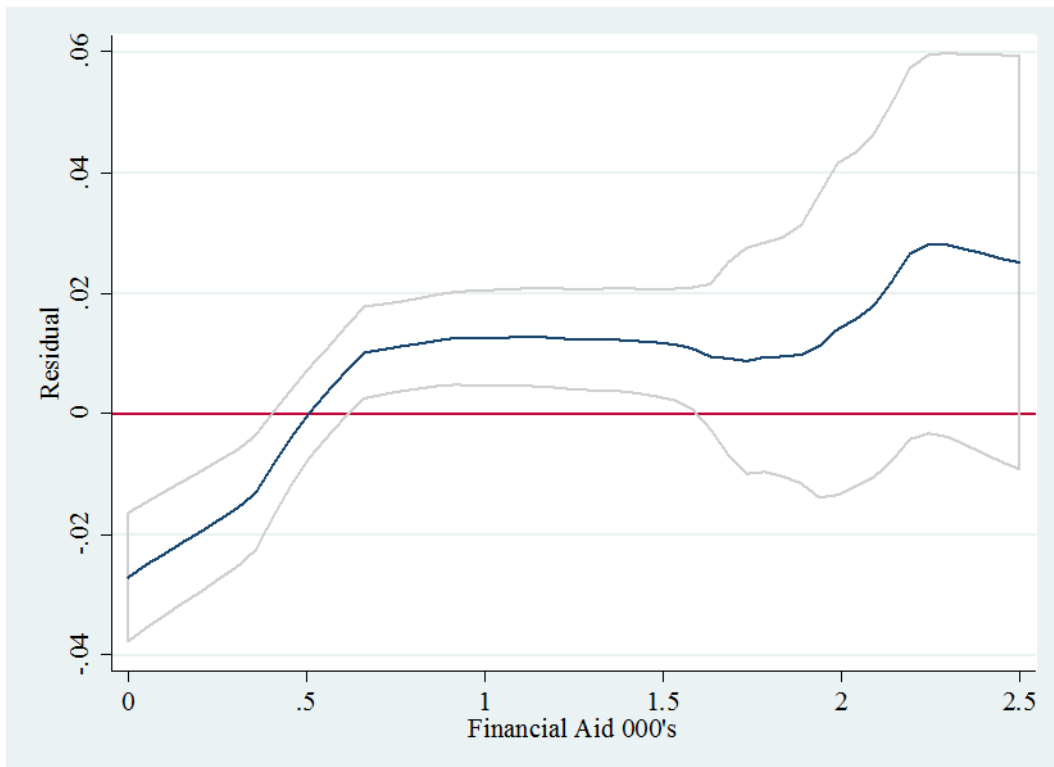
Notes: Figure 4 shows household income and bursary receipt for every first-year student in 2008 for each university. University 3 shows the compliance in 2010 rather than 2008, as that year of entry is not available for that university. The horizontal and vertical lines show the different bursary levels advertised by the university at each income level.

Figure 5: Residuals by distance from cutoff



Notes: Figure 5 shows the residuals from the main specification (equation 1) excluding the bursary parameters, plotted according to distance from cut-off (for household incomes up to £10,000 above and £10,000 below the nearest cut-off) for discrete changes in bursary amounts over £400 (n=3544).

Figure 6: Residuals by financial aid amount



Notes: Figure 6 shows the residuals from the main specification (equation 1) using all students from the main sample, excluding the bursary parameters, plotted according to bursary aid received.

Table 1: Characteristics and student outcomes

	All Students		Balanced Panel of Students	
	(1) Mean	(2) Std Dev	(3) Mean	(4) Std.dev
Household Income	£23,261	£19,476	£23,288	£19,253
Maintenance Grant	£2,059	£1,076	£2,035	£1,055
Bursary (eligible)	£759	£595	£787	£637
Bursary (awarded)	£753	£632	£775	£666
Entry Points	284.38	86.26	278.63	82.94
Male	0.44	0.5	0.43	0.49
White	0.78	0.42	0.78	0.41
Age on Entry	20.13	4.84	20.05	4.79
Complete				
1st Year	0.95	0.22	0.95	0.21
2nd Year	0.92	0.27	0.92	0.27
3rd Year	0.89	0.31	0.9	0.3
Dropout				
1st Year	0.05	0.22	0.05	0.21
2nd Year	0.03	0.17	0.03	0.18
3rd Year	0.01	0.12	0.02	0.14
Standardised Scores Yr1				
1st Year	0	0.99	0.03	0.96
2nd Year	0	0.99	0.03	0.96
3rd Year	0	0.99	0.04	0.95
"Good Degree"	0.6	0.49	0.63	0.48
N	35,879		22,770	

Notes: All Students consists of students from the nine universities undertaking a degree for the years we have available, including continuing students. Students dropping out are recorded as not obtaining a good degree (Good Degree=0). Those continuing and but not yet completed have no measure of good degree. Balanced Panel of Students consists of the subsample of students that theoretically could have completed their course given their entry date, and the data we have available.

Table 2: Impact of financial aid in form of bursary on probability of obtaining a good degree

P(Good Degree)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
1 st Year Bursary Aid	0.012	0.090***	0.043**	0.034**	0.037***	0.035***
Awarded	(0.017)	(0.018)	(0.016)	(0.013)	(0.012)	(0.011)
Panel B						
1 st Year Bursary Aid	-0.039	0.136***	0.124***	0.142***	0.140***	0.128***
	(0.038)	(0.031)	(0.023)	(0.024)	(0.023)	(0.023)
1 st Year Bursary Aid squared	0.022*	-0.017*	-0.030***	-0.041***	-0.039***	-0.037***
	(0.011)	(0.009)	(0.008)	(0.008)	(0.008)	(0.008)
Marginal effect at mean	-0.005	0.110***	0.077***	0.078***	0.079***	0.071***
	(0.024)	(0.019)	(0.013)	(0.013)	(0.013)	(0.012)
R-squared	0.001	0.024	0.120	0.133	0.133	0.144
Panel C – Reduced form						
Marginal effects 1 st Year Bursary Rules	-0.048**	0.062**	0.044**	0.032**	0.031***	0.032***
	(0.020)	(0.030)	(0.016)	(0.012)	(0.010)	(0.010)
Panel D – Instrumented						
1 st Year Bursary Aid	-0.137***	0.056	0.058**	0.062***	0.059***	0.061***
	(0.037)	(0.050)	(0.029)	(0.023)	(0.020)	(0.019)
1 st Year Bursary Aid Squared	0.052***	0.007	-0.011	-0.017***	-0.015**	-0.016***
	(0.014)	(0.015)	(0.010)	(0.007)	(0.006)	(0.006)
Marginal Effects 1 st Year Bursary Aid	-0.056**	0.066**	0.042**	0.035***	0.036***	0.037***
	(0.022)	(0.030)	(0.016)	(0.013)	(0.011)	(0.011)
R-Squared	0.004	0.023	0.119	0.132	0.132	0.143
Sheas's Adj-P R ² Bursary Aid	0.756	0.591	0.530	0.490	0.475	0.472
Sheas's Adj-P R ² Bursary Aid ²	0.796	0.673	0.611	0.567	0.552	0.548
Parental Income		✓	✓	✓	✓	✓
Student Characteristics			✓	✓	✓	✓
University Effects				✓	✓	
Year Effects					✓	
University*Year Effects						✓

Notes: Coefficients in panel C show marginal effect at mean bursary amount. Good degree defined as being equal to 1 for those students obtaining a first class or upper second class degree, and 0 for all other outcomes, including drop out. Sample consists only of those students whose final outcome can be observed. Standard errors are in parenthesis, and are clustered at institution*year level. * p < 0.1. ** p < 0.05. *** p < 0.01, N=22,770 for all regressions

Table 3: Impact of financial aid in form of bursary on probability of completion

P(Complete)	Complete 1 st year	Complete 2 nd year	Complete 3 rd year	Good degree
	(1)	(2)	(3)	(4)
Panel A				
1 st year Bursary Aid Awarded	0.073*** (0.018)	0.085*** (0.017)	0.087*** (0.017)	0.071*** (0.012)
Panel B				
IV- 1 st year Bursary Aid	0.024** (0.011)	0.031** (0.014)	0.033** (0.016)	0.061*** (0.019)
IV- 1 st year Bursary Aid squared	-0.007** (0.003)	-0.009** (0.004)	-0.009* (0.005)	-0.016*** (0.006)
Marginal effect at mean	0.014** (0.006)	0.016** (0.008)	0.019** (0.009)	0.037*** (0.011)
R-Squared	0.127	0.114	0.115	0.143
Parental Income	✓	✓	✓	✓
Student Characteristics	✓	✓	✓	✓
University*Year Effects	✓	✓	✓	✓

Notes: Coefficients presented are of marginal effects at mean bursary amount. Sample consists only of those students whose final outcome can be observed. Standard errors are in parenthesis, and are clustered at institution*year level.* p < 0.1. ** p < 0.05. *** p < 0.01. For all regressions, N=22,770

Table 4: Impact of financial aid on course scores

P(Course Scores)	Course scores 1 st year	Course scores 2 nd year	Course scores 3 rd year	Good degree
	(1)	(2)	(3)	(4)
Panel A				
1 st year Bursary Aid awarded	0.093*** (0.023)	0.056** (0.020)	0.062*** (0.018)	0.071*** (0.012)
Panel B				
IV- 1 st year Bursary Aid	0.123*** (0.026)	0.079*** (0.030)	0.046 (0.040)	0.061*** (0.019)
IV- 1 st year Bursary Aid squared	-0.036*** (0.008)	-0.023** (0.011)	-0.010 (0.013)	-0.016*** (0.006)
Marginal effect at mean	0.064*** (0.015)	0.041*** (0.015)	0.030 (0.021)	0.037*** (0.011)
R-Squared	0.085	0.096	0.085	0.136
Parental Income	✓	✓	✓	✓
Student Characteristics	✓	✓	✓	✓
University*Year Effects	✓	✓	✓	✓

Notes: Coefficients presented are of marginal effects at mean bursary amount. Sample consists only of those students whose final outcome can be observed. Standard errors are in parenthesis, and are clustered at institution*year level.* p < 0.1. ** p < 0.05. *** p < 0.01. For all regressions, N=22,770

Table 5: Robustness Checks

Specification	IV Financial Aid Rules			
	Complete 1 st Year (1)	Course Scores 1 st Year (2)	Good Degree (3)	Std Prior Test Scores (4)
Main Specification	0.014** (0.006)	0.064*** (0.015)	0.037*** (0.011)	0.005 (0.027)
Include Continuing Students	0.011** (0.005)	0.046*** (0.013)	NA	-0.024 (0.029)
Exclude Outlying Entry Score Universities (1, 2, 8)	0.019** (0.009)	0.062*** (0.021)	0.051*** (0.018)	-0.010 (0.029)

Notes: Sample sizes vary by year as students drop out. Sample consists only of those students whose final outcome can be observed. Coefficients presented are of marginal effects at mean bursary amount. Outlying Universities based on the lack of overlap in prior test scores with other universities. Standard errors are in parenthesis, and are clustered at institution*year level. * p < 0.1. ** p < 0.05. *** p < 0.01

Table 6: Alternative polynomial specifications

Specification	IV Financial Aid Rules		
	Complete 1 st Year (1)	Course Scores 1 st Year (2)	Good Degree (3)
<i>Income</i>	0.014** (0.006)	0.075*** (0.018)	0.037*** (0.012)
<i>+Income</i> ²	0.014** (0.006)	0.073*** (0.016)	0.038*** (0.011)
<i>+Income</i> ³	0.014** (0.006)	0.064*** (0.015)	0.037*** (0.011)
<i>+Income</i> ⁴	0.013** (0.006)	0.051*** (0.019)	0.033*** (0.010)
<i>+Income</i> ⁵	0.014** (0.007)	0.055*** (0.019)	0.034*** (0.011)

Notes: Notes: Coefficients presented are of marginal effects at mean bursary amount. Sample consists only of those students whose final outcome can be observed. Standard errors are in parenthesis, and are clustered at institution*year level. * p < 0.1. ** p < 0.05. *** p < 0.01, For all regressions, N=22,770

Table 7: Heterogeneity

Outcome	All	Male	Female	Enter Young (Age<20)	Enter Old (20>=Age <30)	Poor	Rich	Tariff Low	Tariff High	Relatively Low Tariff	Relatively High Tariff
Good Degree	0.031*** (0.010)	0.046** (0.022)	0.041** (0.019)	0.051*** (0.016)	-0.026 (0.023)	0.207** (0.096)	0.065** (0.028)	0.031* (0.016)	0.092*** (0.019)	0.021 (0.013)	0.072*** (0.016)
Test Scores Yr1	0.057*** (0.016)	0.070** (0.032)	0.085*** (0.027)	0.089*** (0.025)	0.010 (0.046)	0.316* (0.171)	0.082 (0.057)	0.051*** (0.018)	0.121*** (0.039)	0.045 (0.041)	0.071** (0.034)
Complete Yr1	0.014** (0.006)	0.020*** (0.007)	0.014 (0.010)	0.018*** (0.007)	-0.021 (0.019)	0.070* (0.042)	0.014 (0.010)	0.013** (0.006)	0.033*** (0.007)	0.008 (0.006)	0.033*** (0.005)
Obs	22,770	9,740	13,030	17,150	5,620	11,385	11,385	9,795	7,733	7,035	10,493

Notes: Notes: Coefficients presented are of marginal effects at mean bursary amount. Sample consists only of those students whose final outcome can be observed. (Relatively) Low Tariff defined by any student under the 50th entry test score percentile (within their institution). The 5242 students with no recorded entry test scores are excluded from the test score heterogeneity. Standard errors are in parenthesis, and are clustered at institution*year level. * p < 0.1. ** p < 0.05. *** p < 0.01.