Testing Means-Tested Aid

Richard Murphy[*] and Gill Wyness[§]

February 2022

Abstract: We estimate the effect of an unexpected institutional financial aid (IFA) award on student outcomes using administrative data collected from nine universities, exploiting variation in IFA schedules within and across university-entry cohorts. Each £1,000 of IFA during the first year of college increases the chances of completing that year by 1.4 percentage points, improves test scores by 0.059σ and increases the chances of graduating with a good-degree by 3.4 percentage points. We find high ability and low-income students benefit the most and calculate outcome maximising, and cost minimising IFA schedules for each university.

Acknowledgements: We thank Damon Clark, Susan Dynarski, Caroline Hoxby, Sarah Turner, John Van Reenen and Etienne Vouga 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.

Correspondence: g.wyness@ucl.ac.uk; richard.murphy@austin.utexas.edu
* University of Texas at Austin, Population Research Center, NBER and Centre for Economic Performance, London School of Economics
§ UCL Institute of Education and Centre for Economic Performance, London School of Economics

1. Introduction

Whilst college participation has grown substantially among young people from the poorest backgrounds over the past 20 years, their completion rates have remained stubbornly low (Bailey and Dynarski, 2011). Moreover, low-income students are more likely to drop out of college or perform poorly in courses even conditional on prior attainment (Crawford et al, 2016). Given the substantial returns to undergraduate degree completion and degree class (Card, 1999; Feng and Graetz, 2017), improving the degree performance of disadvantaged students is a key challenge for education policymakers. This raises the question: can financial aid be used to help equalize the outcomes of students once they are enrolled in college?

Establishing that financial aid helps students to complete college ("the intensive margin"), separate from any enrolment effects ("the extensive margin"), helps us to provide important evidence of the role of credit constraints in higher education. The large literature estimating the effects of financial aid on enrolment (Dynarksi, 2002; Van der Klaauw, 2002; Cornwell et al 2006; Dynarksi, 2008; Cohodes & Goodman, 2014; Castleman & Long, 2016) do not provide direct evidence that credit constraints exist, as any financial aid would also be reducing the net price and would increase the return to investing in college regardless of credit constraints. In contrast, a positive impact of financial aid on student performance, conditional on enrolment, is hard to rationalize without credit constraints playing a role. This is even more evident for students who have enrolled at the same university, which is indicative of similar tastes for schooling (non-pecuniary benefits) and returns to college attendance (Lochner and Monge-Naranjo, 2011). Our paper therefore provides important new evidence of the role of credit constraints in higher education (Lochner and Monge-Naranjo, 2012; Belley and Lochner, 2007; Lovenheim, 2011; Cowan, 2016) and in particular their role in college performance (Denning, 2019; Stinebrickner and Stinebrickner 2008).

In addition, establishing the impact of aid on the intensive margin has important practical consequences for how universities should allocate their resources. If only the extensive margin mattered (e.g., getting students into college), and aid did not help students to progress whilst in university, then universities should target resources towards outreach (OFFA, 2013) and other policies earlier in the education life-course to help ensure young people are academically prepared for college (Caucutt & Lochner, 2020). On the other hand, if aid does help student degree performance on the intensive margin, then institutions could take on students who might appear marginal and then use aid packages to assist them through college. Knowing which types of students benefit from aid is also important, since resources are limited and institutions face the choice of targeting aid on the basis of financial need, academic potential, or some combination.

In this paper we estimate the impact of an unexpected Institutional Financial Aid (IFA) award on a range of measures of college success, where our population of interest is enrolled students, defined by being on universities' rolls (i.e. appearing in their administrative data). There are a number of major challenges in estimating such intensive margin effects. First, is the difficulty of separating the intensive and extensive margin effects. Aid programs are typically promoted by universities to attract students, either away from other universities or from the non-college going population, thereby attracting students who have higher (perceived) gains from aid (Dynarski, 2002). The enrolment of such students, however, would change the composition of the student population towards those individuals less likely to persist (Barr, 2019). A second major challenge concerns the endogeneity of aid: the eligibility criteria of aid programs, either means-tested or merit based, will be correlated with student outcomes (Page and Scott-Clayton, 2016). Third, students have to actively apply, meaning that any estimates are based on the more motivated and informed students (Scott-Clayton & Dynarksi, 2007; Bettinger et al, 2012; Marx and Turner, 2018).

To deal with these three sources of potential bias we estimate the impacts of an IFA scheme which is; 1) not anticipated by the student; 2) has aid schedules whereby equivalent students are eligible for different amounts of aid; and 3) the eligibility and payments occur without student involvement. The IFA we study is known as the English higher education bursary system. It was introduced in 2006 as part of a national increase in tuition fees and required all universities to spend at least 10 percent of tuition fees on means-tested financial aid targeting the poorest students. Despite its size, the bursary program was not advertised to students in university prospectuses or literature, and students do not know which university they will attend with certainty until six weeks prior to enrolment, so there was no means by which students could know what IFA schedule would be applicable to them. We argue that these features make the amount of aid students would be eligible for irrelevant to their college-going decision (Corver, 2010; Callender and Wilkinson, 2013), and test that this form of IFA does not affect the extensive margin of students' decisions to enrol at a particular university. In addition, students do not actively apply for this form of student aid, instead aid eligibility is determined automatically as part of the university application process and is awarded by default. Thus, we can rule out student motivation as a potential mechanism for our estimates.

Bursaries are required to be means-tested and available for the poorest students (defined as those in receipt of the full government-provided maintenance grant). Universities were given complete autonomy in how they designed their bursary systems. This resulted in universities providing different amounts of aid for observationally equivalent students. Importantly for this paper, many universities designed their schemes to have cut-offs or kinks in the slopes in the amount of aid awarded according to parental income, resulting in non-linear variation in the amount received within an institution-entry cohort. We use variation in these schedules in and across universities to estimate the impact of IFA on university success, as measured by retention, annual course scores, and final degree classification.

The goal of our specification is to ensure that our potential outcomes are a smooth function of parental income. In our preferred model, we allow for university-cohort effects as well as flexible and continuous income functions for each entry cohort. Accordingly, identification comes from variation in the IFA schedules across schools (accounting for any level differences) and the jumps or changes in slopes of the aid functions within university-cohort.

To estimate the impact of IFA we collected administrative data from nine English universities over five entry cohorts. These data contain detailed student level information on financial aid, parental income, demographics and prior qualifications, and performance throughout college. Our main sample of analysis focuses on students that entered college between 2006/7 (the year bursaries were first introduced) and 2009/10 (the entry cohorts for whom we have complete academic records).

IFA is paid to students at the beginning of each academic year. Our main specification focuses on IFA received during the first year and outcomes that occur in first year. We find that for the average student, receiving an additional £1,000 of IFA in first year increases their chances of not dropping out of the first year by 1.4 percentage points, and increases first year course scores by 0.059 standard deviations (σ). Ultimately, students who were awarded an additional £1,000 during the first year were 3.4 percentage points more likely to obtain a 'good-degree', defined by graduating with one of the top two degree classes (First, or Upper Second Class Honours) a key determinant of future labour market success (Walker and Zhu, 2013; Naylor et al., 2016; Feng and Graetz, 2017). These results suggest that means-based aid is improving the outcomes of infra-marginal students (Denning, 2019) – those whose initial enrolment was not affected by the IFA but would have worse outcomes without it.

We show that the benefits of IFA are heterogeneous, with low-income and high ability students benefiting most. Both types of heterogeneity are consistent with the presence of credit constraints (Lochner and Monge-Naranjo, 2012). Low-income students have an inability to pay direct costs, and higher ability students will have the highest returns and so larger desire to consumption smooth.

Finally, and similar to more structural approaches found in the literature (Keane and Wolpin, 2001; Johnson, 2013; Abbot et al., 2019), we use our causal estimates to calculate optimal aid policies for each institution in our sample. Our counterfactual analysis for each university is based on the income profile of its students and its current aid budget, either maximising student performance given current spending, or minimising costs while maintaining aggregate student performance. We find that most of the universities in our sample have close to optimal spending allocations. However, there is still evidence of misallocation of aid, with universities typically still not providing the poorest students with sufficient aid, instead providing higher than optimal aid to middle income students.

While a limited number of studies have examined the impact of aid on enrolled students, most are unable to disentangle effects on the extensive margin (e.g. Dynarski, 2003; Bettinger, 2004; Castleman & Long, 2013, Cohodes and Goodman, 2014; Bettinger, 2015). We therefore add to the growing body of work examining the impact of unconditional aid purely on the outcomes of students once they are enrolled (Barr, 2015; Goldrick-Rab et al., 2016; Anderson and Goldrick-Rabb, 2018; Denning, 2019; Barr et al., 2019; Denning et al., 2018).[1] This literature typically finds a positive impact of aid on persistence and graduation.[2] Aside from this, our study makes a number of important contributions.

We are the first to examine the effect of an unexpected aid package that the student has not applied for. The previous studies examining the impact of aid on the intensive margin exploit aid packages which students had to apply for through the FAFSA system (e.g. Denning, 2019; Denning et al., 2018) which is widely understood to be complex

[1] There are also several studies which look at merit-based incentives on the outcomes of enrolled students (Scott-Clayton and Zafar, 2019; Scott-Clayton, 2011; Garibaldi et al., 2012; Joensen, 2013), which typically find such incentives improve student outcomes.

[2] Nguyen et al (2018) present a meta-analysis of studies of aid on enrolment, differentiating these studies according to these different margins.

(Scott-Clayton & Dynarski, 2007; Bettinger et al, 2012), or actively "opt-in" to (Goldrick-Rabb et al., 2016; and Anderson and Goldrick-Rab 2018), meaning that estimates are likely to be identified from the more motivated and informed students. In contrast, the bursary aid program is applied by default and may be unexpected by students. Thus we are able to estimate the impact of aid for students regardless of how well informed or motivated they are, which is important from a policy perspective.

This paper is also the first to examine heterogeneous effects of aid on the intensive margin according to parental income and student ability. This is because the plausibly exogenous variation in IFA occurs at different levels of parental income and student ability, in contrast with much of the existing reduced form literature which typically exploits a single discontinuity or treatment at single point (Goldrick-Rab et al., 2016; Denning et al., 2018, Barr 2015, Denning 2019). We argue therefore that our results are applicable to a broader range of students.

A further contribution is that we use this heterogeneity to provide counterfactual analysis of IFAs at each institution. We see this as a bridge between the reduced form approach and the structural modelling approach, the latter of which aims to draw conclusions about financial aid policies more generally (Keane and Wolpin, 2001; Johnson, 2013; Abbot et al, 2019). However, these papers require many upfront assumptions about the parameters, including preferences and technology, whereas we estimate the causal impact of aid parameters and then optimise. Despite our different approach, we draw similar conclusions, that credit constraints do exist, and that targeted needs or merit-based aid policies are more effective than blanket approaches.

There are several channels through which an unconditional cash transfer might have an impact on student outcomes, such as gift exchange or psychological benefits (Goldrick-Rab et al, 2016, DesJardins et al, 2010).[3] However, the heterogenous nature of our results point to the primary channel being through relieving credit constraints (Belley and Lochner, 2007; Cowan, 2016; Lochner and Monge-Naranjo, 2012, Lovenheim, 2011). Our finding of the continued existence of credit constraints in a setting where students already received loans for the full amount of tuition, in addition to generous amounts of financial support for living costs from national aid programs, speaks to the 'free college' debate and suggests that free college (even coupled with finance for living costs) may not be sufficient to alleviate credit constraints for all students.

[3] Our administrative data does not contain information that would allow us to test for these mechanisms, though our findings suggest that gift exchange is an unlikely mechanism, since it is unclear why poorer or higher ability students would have larger marginal benefits to IFA.

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 Setting

2.1 Financial Aid

The UK higher education system is characterised by large-scale state programs to reduce credit constraints at the point of enrolment, which take the form of deferred tuition fees (repaid through income contingent loans) and financial support to contribute towards living expenses (Appendix B, and Murphy et al., 2019). During the period we study, financial support for living expenses consisted of means-tested maintenance grants and loans. These amount to a considerable 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. All these policies were widely publicised and so would likely impact both the intensive and extensive margins.

The financial aid program that is the focus of this paper is the English higher education bursary scheme, an Institutional Financial Aid (IFA) scheme which works alongside the national aid programs. It was introduced by the UK government in 2006, and required institutions to offer a bursary to all students whose parental income was below £17,500 per year. Apart from this stipulation and the minimum amount of £300 per year for the lowest income students (see Appendix B.1 for details), universities were given complete independence in how much to give out and to whom. Corver (2010) states that "Institutions determine their bursary schemes to meet their own objectives, leading to a range of bursary levels and eligibility criteria across institutions."

In practice, universities typically offered more aid than the minimum required, and extended it to more students than required. Thus, the bursary system became a substantial provider of aid, allocating over £300m annually across England. Around 44 percent of students receive a bursary, with the average amount received around £750 per bursary holder per year. This is around 11 percent of the total aid (of maintenance grants, loans and bursaries) for the lowest income students (See Appendix B).

2.2 Variation of interest

The decentralised nature of this IFA programme provides us with two useful sources of variation. First, the IFA schedules vary non-smoothly within institution by parental income. Second, unlike national financial aid schemes, where it is difficult to find a direct counterfactual for a given student, as students with the same characteristics receive the same amount of aid, the amount of IFA a given student is eligible to receive varies across institutions and entry cohorts. Therefore, our data contain a range of potential counterfactuals at different levels of parental income. This allows us to explore the heterogeneity of the impact of aid by student characteristics.

These features are illustrated by Figure 1, which shows the variation in bursary awarding rules by parental income for each of our nine universities. We see a wide range of IFA values for students from similar income backgrounds. For example, students with zero parental income could receive as little as £350 and as much as £3,200 per year, depending on the university attended and year of entry. In addition to the cross-university variation, there is cross-cohort variation within institution as universities experimented with their schemes across five entry cohorts.[4] Take University 4, 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 at this university decreased from three to two.

Our estimation strategy involves exploiting variation in bursary scheme formulae. Many of the universities have sharp changes in bursaries awarded for a small change in parental income within a university-cohort (universities 1, 2, 4, 6, 7 and 9). For example, in University 4, 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 £15,001 would receive a bursary of £1,545. Others have smoother relationships between parental income and aid but differ in the IFA gradient (Figure 1, universities 3, 5 and 8). We exploit all variation arising from the bursary schedules to identify the impact of received aid conditional on a smooth and flexible parental income function.

[4] Note, whilst our period covers university entry between 2006 and 2011, some of our universities are left or right truncated.

2.3 Extensive Margin

We can identify the impact of IFA on the intensive margin because the institutional setting makes it extremely difficult for students to know which bursary scheme would be applicable to them. This uncertainty stems from the centralised nature of the UK university application process, which incorporates the student aid application. During the autumn of the final year of secondary school, students can apply to up to five colleges through the UCAS system (see Appendix B3 for details). During this application, 92 percent of students state their parental income (Bolton, 2016; SLC, 2013) in order to be considered for the national maintenance grant and student loan program. This information is used (and validated) by the central government to administer these programs. Importantly, this information is also given to the universities for them to administer their own IFA. This information is only provided after the application process. Students then receive offers from universities, conditional on secondary school qualifications. They choose their preferred university, and are contractually obliged to enrol on that course if they meet the entry requirements (unless they formally withdraw their application for all universities that year, see Appendix B3). Months later, students sit their exams, the results of which are received in mid-August and which determine which university they will attend, beginning late September.

This centralised process leads to four institutional properties that make the IFA scheme non-salient (See Appendix B.2 for details relating to each):

i) IFA schemes are not advertised by universities, and there is no centralized comparison tool to help students understand what is on offer.

ii) Students do not know which university they will attend, and therefore which bursary scheme would be applicable, until a month before courses begin.

iii) Students do not apply for IFA; their eligibility is automatically calculated during the centralised application stage if they provide information on parental income, which is required to assess their eligibility for the generous state aid programs.

iv) Students are not notified of aid offers from universities.

Aside from these institutional features, two studies provide empirical support that bursaries do not impact the enrolment choices of students in England. Callender and Wilkinson (2013) survey students who enter English universities in 2008, coinciding with the mid-point year of our sample period. They find 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."* They continue, *"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."* In addition, Corver (2010) finds no evidence that the bursary a student would be eligible for impacts their rank preferences for universities among their offer set.

Given the empirical evidence and institutional details, we are confident that bursaries do not, and indeed cannot, impact on the extensive margin. Aside from the centralised system making the IFA schemes less salient, they also reduce the possibility of students gaming the system. Students submit their parental income at least seven months before they know which IFA schedule will apply to them, and this parental income is validated by a government department against tax records. The possibility of students sorting to universities on the basis of IFA, and gaming the system in terms of reported parental income are both tested for our sample in Section 5.1.

3. Data and institutional compliance

3.1 Data

Our sample of interest is enrolled students who appear in an institution's administrative records. We collected administrative data from nine UK universities, capturing the entire undergraduate population of domestic students for up to six cohorts of students who began their studies between 2006 and 2011. To obtain this data, we contacted all 159 higher education institutions in the UK, asking them for individual level student data on attainment, parental income and bursaries awarded for all enrolled students. Of these, we received data from 25 England-based institutions, covering 341,398 students.

As our estimation strategy relies on exploiting discontinuities in institutional financial aid functions which occur for small changes in parental income, we further restricted our sample to nine universities, for whom data on parental income is observed for all students. We limit our sample to students with parental income below £50,000 as some of our universities did not provide data for students above this point (median household income in the UK was £28,100 in 2010/11 (ONS, 2012)). Note that, with the exception of one cohort (2010) in one university (university 5), no bursaries were provided to students above this level of parental income. We then discarded students undertaking vocational courses or those above or below a bachelor's degree level (2.4 percent) as they are not eligible for the IFA we study. These three restrictions reduced our sample to 35,879 students.

Appendix Table A1 shows mean student characteristics at our universities compared to all other higher education institutions in the UK. Our universities are representative of the sector as a whole in terms of gender, ethnicity and disability. Students in our sample are slightly younger (42 percent are under 21, versus 37 percent of students in the sector as a whole), and are more likely to get a first or upper-second-class degree than the UK average. Appendix Figure A1 shows where our universities fall in the distribution of all other UK universities according to the following student characteristics: gender, age, disability, domicile, ethnicity, and entry test scores. The universities in our sample are found throughout the distribution on all of these metrics, indicating that our sample is representative of the UK sector as a whole.

We observe four cohorts of students all the way through their studies (3 years), whilst we can only observe the first or second year of students who started later. Our preferred specifications use only the sample of 22,770 students for whom we observe for three years (entering cohorts 2006/7-2009/10). In a robustness check, we estimate the impact on outcomes including continuing students (i.e. the full sample of 35,879 students). The dataset tracks students throughout the course of their three-year degree, including whether they dropped out, their year of drop out, their average annual course scores, and their final degree classification.

Table 1 presents descriptive statistics on the individuals in our full and main samples. The average IFA received per student during the first year is £775, and 76 percent of the sample received a non-zero amount. To compare this to the state aid programme, in 2006/7 a student with zero reported parental income would receive a total of £6,255 comprising of maintenance loans (£3,555) and grants (£2,700). In our sample the average IFA that this student would receive in the first year is £1,138, which is an additional 42 percent of grant aid and 18 percent of all maintenance aid consisting of grants and loans (12 percent if including tuition fee loans in the aid total).

Our first major outcome of interest is first year completion rate. In our sample the overall university completion rate is 90 percent. This is comparable to the national completion rate of around 92 percent (HEFCE, 2013). Dropout is highest in first year, at 5 percent, and steadily declines. Note that the UK has particularly high completion compared to other European countries (Grove, 2014) and has a low transfer rate of students between colleges. Over our period of analysis just 2.3 percent of traditional age first year students transferred to another institution (HESA, 2010; 2015). This is likely in-part due to the inability of students to transfer course credits across universities, as 'credit transfer is not widely accepted in the UK and there is little evidence of institutional practice' (Pollard et al, 2017). Since the vast majority of students who transfer have to restart their degree from scratch, when students do not complete in our sample, we interpret and refer to this as dropping out.

Our second outcome variable, course scores, is the average score awarded to each student across modules each year. At a typical English university, students will take eight modules per year and are awarded scores in each of them. Given the high degree of autonomy across institutions (and within institutions, by course), course scores may not be directly comparable across universities, subjects, or time, so we standardize this measure by subject area, university and cohort.

Our third outcome measure is whether a student obtains a good-degree. Unlike in the US, students in England rarely drop out of college, but this means that 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 use the Naylor et al. (2016) definition of a student obtaining a good-degree as those being awarded a First or Upper-Second Class degree. Students who do not complete the degree are assigned as not achieving a good-degree; therefore, this measure acts as a combination of student performance and dropout rate.

The good-degree outcome is of particular interest as this level is a key differentiator for UK employers. Indeed, graduates with a First or Upper-Second Class degree (also known as a 2.1) have been shown to earn around 8 percent more than those with lower class degrees (Feng & Graetz, 2017; Naylor et al., 2016; Walker & Zhu, 2013). This is also often the minimum requirement for entry to graduate programmes. Of all first-year enrolees in our sample, 63 percent obtain a good-degree by this definition. As a point of comparison, the six-year graduation rate for students who started in the fall of 2006 was 60.5 percent at US public four-year colleges, and 62.5 percent at private non-profit colleges (Shapiro et al., 2014).

3.2 Non-compliance

Despite the strict institutional setup described in Section 2, we observe a degree of non-compliance in our data. This is illustrated in Figure 2, which plots household income against IFA eligibility and IFA receipt, for every individual student for our nine universities in one particular year (2008). The vast majority of students receive the IFA amount that corresponds with their observed household income. However, across all our universities, we observe an average of around 5 percent of students receiving an IFA award that is "too high" and around 7 percent receiving an award that is "too low".[5]

A concern is that this non-compliance is arising from systematic factors that could generate biases if we were to estimate our models using simple OLS. Administrators at these universities stated three situations in which the amount of IFA received does not equal the amount that should be received based on reported parental income.

The first case is when a reassessment of parental income indicates that the student would be eligible for a different student aid amount. This could be due to an administrative error or a change in circumstances. This is unlikely to be due to systematic under-reporting, as the parental income measures are validated by the government. These errors or changes in circumstance would result in classical measurement error, which would bias OLS estimates downward.

The second type of non-compliance concerns student 'pre-dropout'. If a student enrols in a course but then withdraws before physically setting foot on campus, they will not receive an IFA. They would still, however, be recorded in the administrative records, and hence will be in our population of interest of enrolled students. For this type of student, their aid received would be correctly recorded as zero despite being potentially eligible, and their outcome would be correctly recorded as dropped out. However, an OLS estimator would not recover the marginal effect of IFA because, for these students, aid received is endogenous to their 'pre-dropping out.' In other words, the fact that the student has dropped out determines their IFA received to be zero. This is an example of reverse causality, where the students' outcome determines the treatment status. Using an OLS estimator in such a setting would recover a biased estimate of the impact of IFA.[6]

The third example of non-compliance concerns the university using its discretion to award additional funds to some students, which we refer to as implementation error. If institutions systematically award high ability students more than they are entitled to, naive OLS estimation would again bias our estimates upwards. We address these concerns using a simulated instruments approach.

[5] The raw correlation between aid awarded and eligible is 0.886 (Appendix Figure A2).

[6] Of the students that received no IFA despite being eligible, 73 percent have recorded test scores, and 68 percent complete the first year, meaning that they did not pre-drop out. This implies that the majority of the non-compliance where a student receives zero aid but is eligible for aid, is not due to pre-dropout. We cannot determine precisely how many students do pre-drop since these students are observationally identical to students who received zero aid in error, and dropped out in the first year. Such students make up only 1 percent of the main sample (241 students). Appendix Table A3 present robustness estimates excluding this subsample of students.

4. Methodology

4.1 Modelling Student Outcomes

The main analysis focuses on IFA received during the first year of university and first year outcomes. Using the IFA aid received in subsequent years is problematic for two reasons. First, aid received in later years would be endogenous if first year aid impacts dropout rates. Second, aid in subsequent years is based on concurrent parental income, running the risk of students/parents potentially manipulating income (e.g. by declaring a marital separation) based on their learned knowledge of the bursary schedule.

Student performance depends on a students' ability, as well as other personal characteristics, and factors that impact student liquidity: family income, state-based aid (i.e. student loans and grants), and IFA.[7]

We represent this simply in equation (1).

$$y_{ijt} = h(A_i) + f_{jt}(I_i) + \rho X_i + \varepsilon_{ijt} \tag{1}$$

where y_{ijt} is the outcome of student i attending university j, who started in entry cohort t. Student outcomes are a smooth function $h(.)$ of the amount of IFA received by student i in thousands of pounds in their first year at university, A_i. Our goal is to estimate the average marginal impact of institutional financial aid, distinct from other forms of state-based financial aid. We allow for different functions of IFA, $h(.)$, and in each case present the average marginal effect.

Parental income I is accounted for with a smooth functional form $f_{jt}(.)$. This allows for a smooth relationship with parental income to vary by both university attended, j, and entry cohort, t. This accounts for endogenous sorting to universities, and changes to the mapping between outcomes and parental income over time and university attended. With such a function, identification of $h(.)$ would be leveraged by jumps and kinks in the aid functions, and differential gradients across universities.

To ensure that potential student outcomes vary smoothly with parental income we control for student characteristics X_i in a flexible way, using dummies for each age, gender, ethnicity type, and a set of 30 dummies for university entry grades (each being an indicator for a 10-point range in entry grades spanning 0-300). Note that we do not

[7] Tuition fee payments are not included in our simple framework because tuition fees are fully deferred, and so do not act as a credit constraint while enrolled in college (See Appendix B1).

include state-based grant and loan aid in A_i. This is because these forms of state aid are likely to affect the extensive margin as students are well informed about these programs (Deming and Dynarski, 2010). Including state-based aid in A_i would prevent any aid parameter from being a pure intensive margin parameter. The $f_{jt}(.)$ function partially accounts for state aid programs as they are also determined by income. However, state programs will not be fully accounted for as they are not smooth functions of income, as they also contain kinks. We therefore account for them by conditioning on maintenance loans and grants directly and allowing for diminishing returns by including quadratic polynomials for state loans and grants separately in X_i. These values are imputed using the stated parental income and state aid rules for the student's enrolment year.

4.2 Estimating Impact of IFA on degree success

The primary concern with the model above is the possibility that unobserved factors affecting degree success are correlated with college aid receipt. To address this problem, we employ an estimation strategy that takes advantage of jumps and kinks in the bursary aid schedule, as well as variation in schedules across universities, to estimate the effect of aid on outcomes. Our empirical strategy is somewhat similar to Clark and Del Bono's (2016) "heavily parameterized regression design" and the methods employed by Dahl and Lochner (2012) in their study of the impact of family income on child achievement. Central to our analysis is the variation in IFA schedules within schools due to jumps and kinks at different values of parental income, and the variation in IFA schedules across schools.[8]

Key to our estimation strategy is to correctly account for the relationship between parental income – the main determinant of IFA within a university-year – and student outcomes. As set out in equation (1), the ideal specification would estimate separate, fully-flexible income polynomials for each university-cohort $f_{jt}(I)$. In this most general case, the relationship between income and outcomes would be allowed to vary systematically between universities, j, and cohorts, t, thus leveraging only jumps/kinks in the aid functions. This is similar to a control-function approach, since any unobserved factor which is a smooth function of parental income would be accounted for through the flexible income function. Critically, this is dependent on the functional form for income

[8] We do not to use a regression discontinuity approach for three reasons, all relating to maximising power. First, there are many close cut-offs in some university-cohorts meaning that there are few students each side of any particular cut-off. Second, there is no standard change in the amount of aid received, so imposing a uniform treatment would increase the standard error of the treatment parameter. Third, some universities have kinks in addition to jumps in financial aid. Using a smooth, flexible functional form for parental income allows us to efficiently exploit all these sources of variation.

being sufficiently flexible to account for such factors. Using the most general control function $f_{jt}(I_i)$ would achieve this, however that would be very demanding on the data. Recognizing this, we impose some restrictions on $f_{jt}(I)$ in most of our analysis. We explain (and test) these restrictions below.

A simple approach to imposing some structure on this relationship would be the following: $f_{jt}(I) = f(I) + \theta_{jt}$. This allows for level differences in outcomes across universities and cohorts but imposes the assumptions that income has a constant relationship with outcomes across universities and cohorts. There are reasons to question both of these assumptions.

Over our study period, there was increasing emphasis on improving university completion and attainment of students from poorer backgrounds (known as "widening participation" students) (Chowdry et al, 2013, Murphy et al, 2019), and the relationship between state financial aid and parental income changed. Therefore, we are concerned that the relationship between parental income and outcomes would vary across cohorts, $f_t(I)$. In addition, the relationship between parental income and degree performance could vary across universities, $f_j(I)$. This could arise due to differences in the cost of living at different institutions or through differences in university admissions policies, academic programs, and student populations.

The importance of allowing for either or both of these interactions with parental income is one that can be empirically tested. Our main analysis assumes that $f_t(I)$ is the most appropriate specification. Through empirical testing we show that this specification cannot be rejected for more general cases. This provides us with the following:

$$y_{ijt} = h(A_i) + f_t(I) + \rho X_i + \theta_{jt} + \epsilon_{ijt} \tag{2}$$

where $f_t(I)$ is a flexible fourth-order polynomial of parental income interacted with cohort of entry t, and θ_{jt} is a set of university-cohort fixed effects.[9] Thus, we allow for different outcome levels across institutions and over time, and for the relationship between income and the outcome variables to vary across cohorts, but not across institutions.

As described in Section 3 (Figure 2), there is some non-compliance to the IFA schedules. This could be due to i) measurement error, ii) non-compliance on behalf of

[9] Table 4 presents estimates for alternate order polynomials in parental income which are interacted with university, cohort, and university-cohort, and establishes that our results are robust to other functional forms.

the student (the student not attending after being accepted) resulting in reverse causality, or iii) non-compliance on behalf of the university (the university administering the wrong amount of aid). To account for non-compliance of bursary received, we instrument the amount of IFA actually awarded, A_i, with the amount of IFA the student is eligible for, E_i. We determine E using each universities' funding schedule in combination with recorded parental income, to create a simulated instrument. We implement a two-stage least squares approach using this simulated instrument in the first stage where β represents the compliance rate of students and universities to the IFA eligibility rules. The second stage estimates the relationship between students' aid as predicted by the rules and the outcome of interest. We use the following specification:

$$y_{ijt} = h(\widehat{A_i}) + f_t(I_i) + \rho X_i + \theta_{jt} + \varepsilon_{ijt} \tag{3}$$
$$A_i = \beta E_i + \eta_t(I_i) + \gamma X_i + \delta_{jt} + \mathcal{E}_{ijt}$$

To account for any diminishing marginal returns to aid in $h(A_i)$, our preferred specifications will include linear and quadratic terms for IFA. This leaves us with two potentially endogenous variables (aid and aid squared), which we instrument for with eligible aid and its square (Wooldridge, 2009; Dieterle & Snell 2014).[10]

Using eligibility rules to address biases due to measurement error in parental income (as in our first example of non-compliance) is a well-established technique. The use of this simulated instrument also addresses biases due to non-compliance to the rules by the student or the university. To see why, consider estimating the impact of E on student outcomes. This reduced form estimate would not even need to measure IFA received, A. If the rules are not well complied with (e.g. many students who are eligible for aid choose not to enrol and so do not complete the first year), then the resulting intention to treat effect would underestimate the true impact of IFA. A Wald estimate, using E as an instrument for A, would scale this by the compliance rate, producing a local average treatment effect for the compliers.

Using E_i as an instrument for student non-compliance requires that eligible aid correlates with received aid (relevance condition) and that eligible aid does not impact the decision to pre-dropout (exogeneity condition). Because students do not receive

[10] Note, one does not generally get consistent estimates using 2SLS to instrument for higher order polynomials in a mismeasured variable (Amemiya, 1985). However, 2SLS estimates of the linear and quadratic terms are consistent in the case where h(.) is a quadratic function. This requires that eligible aid, our instrument for aid received, satisfies independence and not just uncorrelatedness.

information about their eligible IFA (see Appendix B.2), their decision to not attend could not be due to the rules.

4.3 Post First Year Outcomes

Estimation and the interpretation of the average marginal effects on first year outcomes is straightforward. Estimation of the effects of IFA become more complex on outcomes after the first year for two reasons. First, the receipt of IFA during a subsequent year is dependent on the student completing a previous year, and completion is a function of IFA, meaning that receipt of IFA in the second and third years is a function of all previous IFA received.

Second, the IFA received in the second year is dependent on the IFA schedule of that university and the students' parental income in that year. Because students may become familiar with the IFA schedule, they could then game stated parental income. Our initial estimates for the impact of IFA on obtaining a good degree only use first year IFA, so the marginal effect could be interpreted as being allocated a given treatment status in the first year and the associated correlations with later treatments. Notice that first year IFA would be representative of annual IFA received each year, if neither parental income nor the thresholds change.

We also consider a specification that uses a running average of IFA received up until the relevant point in time $\overline{A}_\iota$, which is instrumented by the running average of potential eligible IFA, $\overline{E}_\iota$. This $\overline{E}_\iota$ is based on the eligibility rules, up to and including the current school year, calculated on the basis of first year income. This accounts for any changes to the IFA formulas that occur over time, and is generated for all students regardless of dropout status, making it independent of previous outcomes.[11] For example, for second year outcomes, the appropriate $\overline{A}_\iota$ is the average of IFA received during the first and second year, and the corresponding $\overline{E}_\iota$ would be the eligible IFA over the first and second years based on first year parental income regardless of student enrolment. Another issue that arises for subsequent years is that students who drop out will not have test scores. Thus, our course score estimates would be biased by sample selection. For test scores in subsequent years, we additionally estimate bounds for the effects of $\overline{A}_\iota$ (Lee, 2009).

[11] Holding parental income constant, there is over a 0.95 correlation of first year eligible IFA with eligible aid amounts in the second and third years, conditional on remaining enrolled. Allowing parental income to vary over time reduces this correlation to 0.83 in the second year and 0.75 by the third year, conditional on remaining enrolled.

5. Empirical results

A key assumption of our approach is that students are not sorting to universities (or university cohorts) on the basis of the IFA they would be eligible for. In Appendix C we present empirical evidence of this by showing no significant changes in densities or pre-college observable characteristics associated with changes in eligible IFA. First, we use a manipulation test proposed by Cattaneo, Jansson, and Ma (2016), which builds on the McCrary (2008) test. We apply this to all students near a discontinuity of over £100 and show that there are no significant differences in student attendance across these boundaries. This is supported by visual evidence for the lack of bunching with Appendix Figure A3. Second, we show that IFA receipt is not predicted by pre-college observable student characteristics conditional on $f_t(I_i)$ and θ_{jt} (Appendix Table A2).

5.1 Graphical Results

Before presenting estimates from our main specifications, we show suggestive graphical evidence of the impact of aid that relies on fewer parametric assumptions. Using the same subsample of students as above, whose parental income is within £10,000 of their nearest IFA cut-off of more than £100 in aid (n=9,093), we estimate specification 2 (omitting the aid parameters) for first year test performance and plot the residuals from this regression against distance from cut-off. This allows us to focus on the variation in outcomes that is unexplained after conditioning on $f_t(I_i)$, θ_{jt} and student observables. Figure 3 presents a simple local linear plot of these residuals each side of the stacked discontinuity. We see that students just below the IFA cut-off who receive a discontinuously larger amount of aid, have better outcomes than those just above the aid threshold.[12]

In Figure 4, we plot residuals from the same specification against eligible aid, this time using the main sample. This semi-parametric method of displaying the residuals has the advantages of not assuming a functional form of the aid function $h(.)$ while flexibly accounting for observable differences across students. Here, the negative intercept implies that students with zero eligible aid on average do worse than the average student conditional on observable characteristics. We see that, over the first £1,250 of bursary

[12] Rather than plotting the residuals, Appendix Figure A4 plots standardized first year test performance for students close to IFA cut-offs (same definition as Figure 3), who are also below median income. Using this raw outcome measure on this subsample, we again see a discontinuous jump in student performance around the cutoff.

aid, students perform better than those with zero aid, but for IFA greater than £1,250, the gradient levels off. For levels of aid above £1,750 confidence intervals become extremely wide, as few individuals receive bursaries at such high levels. Given this observed relationship, we present estimates allowing for decreasing marginal returns to aid in our specifications by including a linear and squared term for financial aid received.

5.2 Main Results

A. First Year Completion, First Year Course Scores and Degree Class

Table 2 reports estimates of the impact of first year IFA on three student outcomes: the probability of completing year 1 (column 1), course scores in year 1 (column 2), and obtaining a good-degree (column 3). We present OLS estimates from specifications 2 and 3 described in Section 4.2, allowing for fourth order parental income polynomials for each entry cohort, university-cohort effects, and student characteristics.

In Panel A we use specification 2 and assume constant returns to financial aid. Here, our findings imply that each £1,000 of aid results in a 3.5 percentage point increase in the probability of a student completing first year, a 0.039σ increase in first year course scores and a 3.6 percentage point increase in the probability of a student earning a good-degree.

In Panel B, we allow for decreasing marginal returns to aid by adding a quadratic term and report the effects for aid and aid squared. For ease of interpretation, we also present the average marginal impact of £1,000 of IFA, which given the quadratic specification is also the marginal effect at the mean IFA (£755). The estimates imply that a £1,000 increase in bursary aid at the mean increases the probability of completion by 7.4 percentage points. This effect is considerably larger than the linear estimate of 3.5 percentage points, implying that diminishing marginal returns set in relatively quickly, as can be inferred from the comparably large negative squared term. Similarly, we find considerably larger effects for improving course scores (0.091σ) and degree class (7.2 percentage points).

To address the potential for non-compliance, we repeat the analysis in Panel B using IFA eligibility, instead of IFA received. This reduced form set of results is presented in Panel C. These parallel results are considerably lower than the estimated effects of aid received for all three outcomes. The average marginal effect for completion of the first year is 1.2 percentage points, 0.05σ for course scores, and 2.9 percentage points for good-degree. Using the eligibility rules provides an intention to treat effect and should be lower than the true causal effect of IFA received since some students received different amounts of aid than they are eligible for. Still, they are valid estimates of IFA eligibility regardless of any non-compliance by the student or university. In Appendix Table A3 we explore non-compliance in more detail, and show that the instrumental variable approach is working as expected.

In Table 2 Panel D, we present our preferred estimates using specification 3, which instrument IFA received and its square with the amount eligible and its square. As expected, these 2SLS estimates are lower than the OLS estimates and are higher than the reduced form eligibility estimates. As before, there are decreasing returns to the amount of IFA, with the maximum impact that aid could have occurring around £2,900 (Panel D columns 1-2). Here, the average marginal effect for a £1,000 increase in IFA implies an increase in the possibility of completion by 1.4 percentage points (significant at the 1 percent level), an increase in course scores of 0.059σ, and a 3.4 percentage point increase in the probability of obtaining a good-degree.[13]

We test for weak instruments using the adjusted F-statistics using the method of Sanderson and Windmeijer (2016) as well as the more traditional adjusted Shea's partial R- squared, which take the intercorrelations among instruments into account. Both indicate that the instruments are very strong predictors of the endogenous variables.

B. Degree Completion and Course Scores in subsequent years

Table 3 explores the impact of receipt of average IFA up until the outcome is realized. Panel A presents the reduced form estimates for the average marginal effect of $\overline{E_\iota}$, derived from the linear and quadratic terms for $\overline{E_\iota}$. Note that, for first year outcomes, the estimates will be identical to the reduced form estimates from Table 2, as both use eligible IFA in the first year. Panel B presents the equivalent marginal effects where we use the running average for eligible aid as an instrument of the running average of aid received $\overline{A_\iota}$.

Each column represents a separate regression and set of results for a different outcome. In columns 1-3 we show the impact of an additional £1,000 eligible IFA each year on completion of the first, second and third years. We find evidence that the impact of $\overline{E_\iota}$ on completion of each grade is approximately stable. The marginal effects of $\overline{A_\iota}$,

[13] In Appendix Table A4 we present a parallel set of results for other classes of degree outcomes in addition to 'good degree' (First or Upper Second Class degree). The change in outcomes are largest for this degree outcome, but the gains from aid are found for each degree class level.

instrumented by $\bar{E}_i$ are larger, as expected. Here we also see the cumulative effect of IFA on completion increasing (insignificantly) from 1.4 percentage points in 1st year to 2.1ppt in year 2, and to 2.7ppt in year 3.

Columns 4, 5 and 6 of Table 3 show the marginal impact of IFA on mean standardised course scores each year. Here, although we see a positive impact of IFA each year, the magnitude decreases each year and is significant only in years 1 and 2. Course score estimates are impacted by sample selection as students who drop out before the examinations (including during the first year) are not included. Therefore, in Appendix Table A.5 we present bounded estimates (Lee, 2009), in which we replace missing values of course scores with the standardized minimum and maximum test scores within each university, course and year of entry group. For the first and second year of study we find lower bound estimates of 0.039 and 0.031 respectively. By the third year, the lower bound is insignificant but remains positive, at 0.001, while the upper bound is both positive and significant at 0.141.

In summary, we find a positive impact of IFA aid on completion in first, second and third years, by as much as 2.7 percentage points, and on improvements in course scores in all three years. These effects culminate in an overall improvement in final degree classification of 4.5 percentage points per £1,000, against a mean good-degree rate of 62 percent (Table 5, Column 7).

The effects for completion (the outcome most studied in the literature) are comparable with those found by Denning (2019) who finds that college seniors are 1.8 percentage points more likely to graduate a year earlier with an additional $1,500 of loans and grants. Our estimates are smaller than those found by Barr (2019), who studies an information campaign that led to small reductions in student loan borrowing ($220), and finds a 3.6 percentage point reduction in likelihood of any enrolment a year after. Our effects are also comparable with evidence from Nguyen et al (2019), whose systematic review of the effects of aid conditional on enrolment showed an increase in yearly completion of 1.8 percentage points for each $1,000.

5.3 Robustness Checks

In our preferred model, we allow for university-cohort effects, and flexible and continuous income functions for each entry cohort. In Table 4 we present estimates of the impact of IFA on first year test scores using varying orders of the polynomial in income, interacted with entry cohorts $f_t(I_i)$ (rows of the Table 4) and interacted with university attended $f_j(I_i)$ (columns of Table 4). In addition, we also include estimates of the marginal impact when allowing for university and cohort interactions with parental income, $f_{jt}(I)$.

We choose our preferred specification based on comparisons of Wald-Statistics. The estimate from our preferred specification can be found in column 1, row 5 (0.059). Regardless of the specification, the average marginal impact of aid on student test scores remains relatively constant between 0.072 and 0.052. The exception is the fifth column, which has a smaller estimated effect. Given that our IFA rules predominantly vary at the university parental income level, this specification, with a quartic in parental income interacted with university attended, may be too demanding on the data. This is reflected in the standard errors increasing by half in column 5.

Below the estimates in columns 2 through 5 we provide the p-value for the Wald test that the additional parameters from the preceding column are jointly equal to zero. The university interaction terms with parental income are not rejected, implying that they are not required. This is reflected in the relative stability of the estimates. All these patterns in the estimates are reflected in the final row which allows for university and cohort of entry interactions. Companion tables for completion of first year and degree class are found in Appendix Table A6.

In Appendix D we perform a series of robustness checks using different samples of students to establish the robustness of our estimates of the average marginal effects: including students who have yet to graduate; excluding universities who are outliers in terms of prior academic achievement of students; excluding universities who are outliers in terms of amount of IFA awarded. For each subsample the average marginal effects are not statistically different from those found using the main sample.

5.4 Heterogeneity

We now consider whether the relationship between IFA and outcomes varies by student characteristics. Estimated effects by secondary school performance and parental income are presented in Table 5, and Appendix Table A7 presents estimates by gender and age group.

There is little difference in the impact of IFA by gender, although the estimated marginal effect on first year test scores is slightly (although not significantly) larger for females. This is in line with findings by Angrist et al. (2006), Page et al. (2019), and

Scrivener et al. (2015), all of whom show that the effects of financial aid on outcomes for men and women are not substantially different.[14]

There are distinct and 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 before they are aged 20), rather than more mature students, on all three outcomes.

Unlike papers centred around a specific discontinuity, we can directly test which types of students benefit from aid after enrolment, in terms of both disadvantage and ability. Looking first at ability, we divide the sample at the median of high school achievement for all students in the sample. We find that, although both high ability and low ability students appear to benefit from IFA, high ability students benefit significantly more in terms of first year completion (3.4 percentage points versus 1.1) and obtaining a good degree (8.5 percentage points versus 2.6).

In terms of family income, we re-run our preferred specification interacting eligible aid and its square with parental income. We then evaluate the impact of IFA at the 25th, 50th, and 75th percentiles of parental income. We find that, for all three outcomes, those from poorer backgrounds benefit more from IFA compared to those from more affluent backgrounds, although these differences are not statistically significant at traditional levels for test scores and obtaining a good degree.

Since any education subsidy would lower the net price of college, observing a positive response of aid on enrolment would not necessarily imply that credit constraints are at play. In contrast, our estimates of a positive impact of aid on completion and test scores, conditional on enrolment, provides more convincing evidence of the existence of credit constraints. It is much harder to see why students who have already enrolled in college would respond positively to aid (in terms of completion and course performance), unless they were credit constrained and having to devote time to employment rather than studying.

The fact that we see greater impacts of £1,000 in aid for low-income students than high-income students also suggests that, despite enjoying a "free college" experience, with no upfront fees, and grants and loans for living costs, students may still face credit constraints in the UK system. One reason for this is that even with sources of aid sufficient to cover tuition, individuals will still want to borrow to smooth consumption

[14] Interestingly, the review by Nguyen et al. (2019) highlights both studies that find bigger effects for men and studies that find bigger effects for women.

(Lochner and Monge-Naranjo, 2011). Moreover, the fact that we see larger responses to IFA for higher ability students (conditional on family income) suggests that these students are additionally credit constrained. This is consistent with Lochner and Monge-Naranjo (2011; 2012) who predict that, since higher ability students generate higher returns on human capital investment, they will want to invest more. The further implication is that university aid packages which target high achieving students may be more effective than those based purely on a means-test. Taken together, these findings highlight the educational attainment benefits of targeting financial aid towards students from lower income backgrounds, especially those who are high achieving.

6. **Designing an IFA Scheme**

6.1 Optimizing an IFA scheme

IFA schedules vary across universities. This may be because universities differ in their student composition, because they have different objective functions, or if the IFA are incorrectly designed. While we cannot know what each university is trying to achieve with their IFA, we can evaluate how well their IFA schedules perform according to plausible objectives given their student populations.

In this section, for each university and each outcome (first year completion rates, first year test scores, obtaining a good degree) we calculate the IFA schedule which would i) maximise average student outcomes subject to current spending, or ii) minimise costs subject to maintaining university level average outcomes. This allows us to compare universities' current IFA schedules to schedules which are the solution to a given optimization. We then evaluate which students universities are implicitly favouring with their IFA schedules.

We consider school decisions problems that weight each individual equally, and consider each outcome for each university separately. Of course, universities may direct money to certain students for reasons that are different from our outcome measures (e.g., to improve their wellbeing at college, or to maximise their longer run outcomes). Therefore, all of the following results should only be considered as optimizations for our specific outcome measures. Finally, we take the student composition of universities as given. The extent to which these schedules can be applied to future cohorts depends on continuing the assumption that potential students are unaware of universities' IFA schedules, and therefore do not affect future enrolment behaviour.

The parameters used come from our preferred specification, conditioning on X_i, θ_{jt}, and $f(.)$ a third order polynomial of parental income. We now generalize the aid function $h(.)$, to allow for differential effects of aid by parental income, interacting the third order polynomial of parental income with a quadratic of eligible IFA, $h(E,I)$[15].

This is estimated on the main sample of all universities for each outcome separately, the IFA parameters from which are then used in a counterfactual policy analysis. Given these parameters we calculate the cost minimizing policy subject to holding average outcomes constant for each university, as well an outcome maximizing policy that holds total expenditure fixed for each university. Both optimisations additionally have the constraint that IFA are always non-negative. This constraint means that the optimisations will equate the marginal benefit of IFA for anyone receiving non-zero aid within a university (See Appendix E for details).[16]

6.2 Graphical Comparison of Current to Optimised IFAs

The cost minimisation and the outcome maximisation will recover a distinct recommended IFA amount for every student depending on the outcome of interest and university. In Figure 5 we plot the values optimizing with regard to first-year course scores against parental income for each university along with the existing aid schedules.

First, with regard to maximising average course scores subject to the current budget (dark grey), all but two of the universities would gain by allocating more IFA to the lowest income students at their institutions, and less to higher income students. University 5 is the starkest example of this, since it actually gives more aid to middle- and upper-income students than lower income students (which appears to be a deliberate intention to offset on a one-to-one basis the reductions in state means-tested aid, Appendix B). On the other hand, universities 4 and 8 provide significantly more to the lowest income students than would be required to maximise average test scores.

Second, we consider the optimal IFA schedules for universities to minimise costs, while maintaining the current aggregate levels of student achievement (light grey). As expected, the cost minimisation schedule is always beneath the outcome maximisation schedule. For some universities (1, 2, 3, 6, 7) the aid schedule under cost minimisation

[15] The estimated coefficients for aid function $h(E,I)$ for each outcome can be found in Appendix Table 8.

[16] Note, that for this exercise we include continuing students so that we can document the changes universities have made to their IFA schedules over time. As seen in Appendix D, Table D1, we do not expect aid to affect these students differently.

is very similar to the schedule with outcome maximisation implying that universities current IFA schedules are not far from either optimization. University 5 again stands out as an institution whose schedule is not well aligned with this objective. The university could maintain the same average course scores with large reductions in its spending on financial aid. Instead of providing up to £2,000 in aid to higher income students, the university could reduce spending to less than £100 aid for all students with parental incomes above £16,000 without reducing overall student test score performance.

In Appendix Figure A5 we consider these maximisation and minimisation problems with regard to the two other outcomes: completion of the first year (Panel A) and obtaining a good degree (Panel B). The schedules are generally similar to those for test scores, suggesting the IFA schedules which maximise test score performance are similar to those which maximise completion. There are, however, some differences of note. First, universities should spend marginally more on the lowest income students if they intend to maximise test scores rather than completion. Second, the optimal IFA schedules with regard to student test scores have steeper negative gradients than those for completion. This would culminate in students receiving zero financial aid at lower levels of parental income if universities focused on test scores rather than completion.

For example, most universities could maximise first year completion rates by capping IFA to students with parental income levels near £40,000. In contrast, most universities could maximise first year test scores, by capping IFA to students with parental income between £34,000 and £38,000. Five of the nine universities in our sample currently have a cap on parental income for IFA receipt between these £34,000 and £40,000 benchmarks.

These patterns are reflected in the heterogeneity by parental income results (Table 5) where the impact of IFA varies more for test scores than for completion rates. That there are differences in how aid should be allocated to maximise different outcomes emphasises that spending needs to be tailored to the outcome universities are most interested in.

In summary, these findings show that if the objective is to maximise average outcomes, most of the universities in our sample (1, 2, 3, 6, 7, 9) should reallocate some of the aid currently targeted toward middle income students towards the lowest income students. That some universities have schedules close to those generated from our optimizations suggests that institutions may be trying to solve these problems. Other universities, however, are allocating spending in very different ways, suggesting that either they are attempting to i) solve similar problems but have bad information, or ii) are solving quite different optimisation problems from those we have set out.

6.3 Numerical Comparison of Current to Optimised IFAs

In Table 6 we examine the differences between current and our estimated IFA schemes in more detail. For simplicity, our discussion of these statistics will focus on first year test scores, using two example universities with similar spending per student (of around £700 per student per year), but with very different spending profiles - University 1 and University 4.[17]

The first row simply contains the mean standardised test scores for each university (standardised prior to sample restrictions). The second row shows the gains from universities instead using the IFA schedule which maximises student test scores while keeping spending constant. For University 1 the gain from changing their system is very limited, improving their student performance on average by 0.5 percent of a standard deviation. In contrast, University 4 would experience a much greater benefit, being able to improve student test scores by 3 percent of a standard deviation without increasing spending.

The third row shows the percentage potential financial savings each university could make from switching to a cost minimising scheme while maintaining current outcomes. Here, if University 1 switched to a schedule which minimised costs, they could reduce spending by 25 percent. In contrast, University 4 could reduce spending by 89 percent.

The final row converts these savings from introducing a cost minimising scheme into savings per student in pounds. Here, University 1, can save £135 per student and still maintain first year test scores. In contrast, University 4 can save on average almost £627 per student by reorganising spending.

By comparing the cost minimised IFA schedules to current schedules, we can determine what "additional valuation" universities are placing on students, over and above the value in terms of the respective outcome. This could be due to the university targeting another outcome, because they wanted to attract these students to the university, or to assuage government regulators. This additional valuation is calculated by subtracting cost-minimised IFA from current IFA eligibility for each student. The

[17] The equivalent statistics for completion of the first year and obtaining a good degree can be found in Appendix Table A9.

difference is the monetary value the school places on the student in addition to the outcome, which is comparable across universities.

For each student we plot these "additional values" against parental income for each outcome in Appendix Figure A6. With the exception of universities 4 and 8, there is a similar pattern in terms of "additional value" across our universities, in which they are placing zero additional value on the lowest income students, but as we move further up the income distribution universities start awarding students more IFA than is required to maintain current outcomes. At some point the additional payments start to decline and by parental income levels of £40,000, most universities are placing no additional valuation on these students. This pattern is consistent with universities wanting to attract (or simply valuing) middle income students.

It might be reasonable to assume that universities did not (initially) know how their students would respond to aid, but may have learned 'what works' over time, and adjusted their aid policies each year in an attempt to arrive at an optimal scheme in terms of cost minimisation. Therefore, in Appendix Figure A7 Panel A we plot the percentage of the current budget that could be saved each year for each university from moving to a cost minimisation scheme.[18] This shows relatively little change in the efficiency of the budget use of universities over time.

An alternate way to measure misallocation that incorporates the amount of misallocated funds per student defines university misallocation each year as the average absolute difference between the university's current schedule and a cost minimising schedule (Appendix Figure A7 Panel B). This measures misallocation in absolute terms rather than as proportion of current budget. Here a value of 1 represents a university misallocating £1,000 per student on average in a given year. Using this metric, universities 4 and 8 misallocate much higher amounts than the other universities. University 4 however, experiences a dramatic reduction in misallocation between 2007 and 2008, which coincided with a reduction in spending on low-income students (Figure 1). While we cannot know if University 4 reacted to their own observations of the success (or lack of) of their IFA spending in this year or are merely mimicking the spending profiles of other universities, they have reduced their misallocation per student by 68 percent. Appendix Figures A8 and A9 repeat these two panels for the outcomes complete the first year and obtaining a good degree respectively, which show similar patterns.

[18] An analogous set of results for outcome maximisation given current budget produces qualitatively similar results available upon request.

Finally, we use our estimates of the heterogenous gains from IFA with respect to parental income used in these simulations to calculate how much of the parental-income test-score gap is closed by the current IFA system, and could be closed by an optimal IFA system. For simplicity we define students as high-income if their parental income is above the median in the sample, and low-income otherwise.

The average first year course scores of low-income students is -0.023σ, and 0.076 for high-income students (Table 5), providing an average gap of 0.1σ. Moving from a system with no IFA to the current IFA system, low-income students on average have benefited by 0.087σ, while high income students have benefited by 0.015σ, meaning the current gap would be 0.07σ larger without IFA. If instead universities were to use an IFA schedule which maximised overall first year performance, low-income students would experience larger gains (0.099σ) and high-income students would experience lower gains (0.007σ), meaning that the income achievement gap would be further reduced. With regards to completion, the gap would be approximately 3 times larger without current IFA. If maximising completion, the gap would be reduced by 2.3 percentage points. [19]

7. Conclusion

This paper examines the causal intensive margin effect of institutional financial aid (IFA) on student performance at university. To do so, we use administrative data collected from nine universities, exploiting non-linear variation in the amount of aid received within an institution-entry cohort. We find that each £1,000 of aid awarded at the beginning of the first-year increases students' likelihood of completing first year by 1.4 percentage points, improves their annual test scores by 0.059σ, and increases their chances of graduating with a good-degree by 3.4 percentage points.

We use these parameters to empirically examine the effectiveness of IFA policies on our sample. Our results reveal that universities' spending behaviour is largely on track, but, depending on the outcome in question, performance could still be improved if more funds were targeted towards the lowest income students. Of course, this is true only with respect to the set of outcomes we study in our optimization. An important message is that when designing an aid package, universities should carefully consider their goals. We find that optimal spending profiles differ depending on whether universities are targeting

[19] High-income students in our sample have a one percentage point higher completion rate (0.948 and 0.958). Without the current IFA low-income students would be 2.2 percentage points less likely to complete, whereas high-income students would be relatively unscathed (0.07 percentage points).

degree completion or student attainment, suggesting one spending profile does not fit all goals.

This paper provides two sources of evidence for the role of credit constraints in student performance. First, it is unlikely that the positive effects we observe are due to a price effect: while IFA does lower the net cost of college, this would impact the extensive margin – it is harder to see how a price effect could result in students doing better at school, conditional on enrolment. Second, our unique setting allows us to show that the benefits of aid are heterogeneous over a range of parental income levels, with low-income students benefiting more than high income students.

This plausibly exogenous variation in IFA at different levels of parental income means that we contribute to much of existing reduced form literature which typically exploits a single discontinuity or treatment at single point. This variation allows us to evaluate the IFA schedules at existing institutions by providing counterfactual analysis for a wide range of parental income levels. This is similar in nature to the structural literature, (Keane and Wolpin, 2001; Johnson, 2013; Abbot et al., 2019) which explores how to cost effectively implement financial aid policies.

There are several ways in which our studies diverge, however. First, these papers all explicitly take into account the role of parental transfers and employment whilst at university, whereas these factors are implicitly encapsulated by our reduced form approach. Second, they are able to provide counterfactual analysis incorporating both grants and loans, while we only consider the role of institutional grants. Third, these papers consider the role of credit constraints on university enrolment whereas our focus is purely on the intensive margin.

Our approach has a number of advantages, however. First, these papers require many upfront assumptions about the parameters, including preferences and technology whereas we estimate causal parameters and then optimise. Second, these papers focus on degree completion whereas we use more detailed measures of performance which impact future labour market performance (Walker and Zhu, 2013). Despite these different approaches, we draw similar conclusions, that credit constraints do exist, and, targeted needs or merit-based aid policies are more effective than blanket approaches (Abbot et al, 2019; Johnson, 2013).

Our results show that credit constraints remain an impediment to student performance even within a higher educational system that provides generous levels of credit to students. In England, during the period we study, all new domestic students automatically qualified for a subsidized loan to cover the full amount of tuition fees (max £3,000), means-tested maintenance loans (max £3,420 in 2008/9), and maintenance grants (max £2,835 in 2008/9) each year they were enrolled. Thus, the fact that we find a positive impact of IFA despite this level of state support speaks to the "free college" debate (Murphy et al, 2019), suggesting that free college, even when supported with aid for living costs, may not be sufficient to alleviate credit constraints in some groups. [20]

Whilst our findings are 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, consistent with Lochner and Monge-Naranjo (2011). 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. The bursary scheme we study directed financial aid towards students purely on the basis of parental income and therefore constrained potential gains. Our findings highlight the likely benefits of an even more targeted approach.

[20] Stinebrickner and Stinebrickner (2008) also identify a group of credit constrained students in a "free college" set up.

References

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L. Violante. 2019. Education policy and intergenerational transfers in equilibrium. *Journal of Political Economy* 127, 6: 2569-2624.

Amemiya, Yasuo. 1985 Instrumental variable estimator for the nonlinear errors-in-variables model. *Journal of Econometrics* 28(3): 273-289.

Angrist, Joshua, Daniel Lang, and Philip Oreopoulos. 2006. Lead them to water and pay them to drink: An experiment with services and incentives for college achievement. *NBER Working Paper* No. w12790

Anderson, Drew M., and Sara Goldrick-Rab. 2018. Aid after enrollment: Impacts of a statewide grant program at public two-year colleges. *Economics of Education Review* 6: 148-157.

Bailey, Martha. and Dynarski, Susan. 2011. Inequality in Postsecondary Attainment, in Duncan, J. and R. Murnane (eds.) *Whither Opportunity*, Russell Sage.

Barr, Andrew. 2015. From the Battlefield to the Schoolyard: The Short-term Impact of the Post-9/11 GI Bill, *Journal of Human Resources*, 50(2), 580-613.

Barr, Andrew. 2019. Fighting for Education: Veterans and Financial Aid, *Journal of Labor Economics*, 37(2), 509-544.

Barr, Andrew, Kelli A. Bird, and Benjamin L. Castleman. 2021. The effect of reduced student loan borrowing on academic performance and default: Evidence from a loan counseling experiment. *Journal of Public Economics* 202: 104493.

Belley, Philippe, and Lance Lochner. 2007. The changing role of family income and ability in determining educational achievement. *Journal of Human capital* 1(1): 37-89.

Bettinger, Eric. 2004. 'How Financial Aid Affects Persistence', In C. Hoxby (Ed.), *College choices: The economics of where to go, when to go, and how to pay for it* (207–238). Chicago: University of Chicago Press

Bettinger, Eric. 2015. Need-Based Aid and College Persistence: The Effects of the Ohio College Opportunity Grant. *Educational Evaluation and Policy Analysis*, 37

Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2012. The role of application assistance and information in college decisions: Results from the H&R Block FAFSA experiment. *The Quarterly Journal of Economics* 127 (3): 1205-1242.

Bolton, Paul. 2016. 'Student Loan Statistics' *House of Commons Library*, Briefing Paper, No1079

Callender, Claire & David 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, David. 1999, 'The causal effect of education on earnings', in Orley Ashenfelter and David Card, (eds), *Handbook of Labor Economics*, vol. 3A, North-Holland.

Castleman, Benjamin L., and Bridget Terry Long. 2016. Looking beyond enrollment: The causal effect of need-based grants on college access, persistence, and graduation. *Journal of Labor Economics* 34 (4): 1023-1073.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma. 2018. Manipulation testing based on density discontinuity. *The Stata Journal* 18(1): 234-261.

Caucutt, Elizabeth M., and Lance Lochner. 2020 Early and late human capital investments, borrowing constraints, and the family. *Journal of Political Economy* 128(3): 1065-1147.

Chowdry, Haroon., Crawford, Claire., Dearden, Lorraine., Goodman, Alissa., & Anna Vignoles. 2013. Widening participation in higher education: analysis using linked administrative data. *Journal of the Royal Statistical Society*: Series A (Statistics in Society), 176(2), 431-457.

Clark, Damon, and Emilia Del Bono. 2016. The long-run effects of attending an elite school: evidence from the United Kingdom. *American Economic Journal: Applied Economics* 8(1): 150-76.

Cohodes, Sarah R., and Joshua S. Goodman. 2014. Merit aid, college quality, and college completion: Massachusetts' Adams scholarship as an in-kind subsidy. *American Economic Journal: Applied Economics* 6(4): 251-85.

Corver, Mark. 2010. 'Have bursaries influenced choices between universities?', Research Paper 2010/06, *Office for Fair Access*, UK

Crawford, Claire., Gregg, Paul., Macmillan, Lindsey., Vignoles, A and Gill Wyness. 2016. 'Higher education, career opportunities, and intergenerational inequality', *Oxford Review of Economic Policy*, 32(4): 553–575

Cornwell, Christopher, David B. Mustard, and Deepa J. Sridhar. 2006. The enrollment effects of merit-based financial aid: Evidence from Georgia's HOPE program. *Journal of Labor Economics* 24(4): 761-786.

Cowan, Benjamin 2016. Testing for educational credit constraints using heterogeneity in individual time preferences. *Journal of Labor Economics* 34(2): 363-402.

Dahl, Gordon B., and Lance Lochner. 2012 The impact of family income on child achievement: Evidence from the earned income tax credit. *American Economic Review* 102(5): 1927-56.

Deming, David., & Susan Dynarski. 2010. 10. College Aid (pp. 283-302). University of Chicago Press.

Denning, Jeffery.T, 2019. 'Born Under a Lucky Star: Financial Aid, Credit Constraints, and College Completion'. *Journal of Human Resources*, 54(3):760-784

Denning, Jeffrey T., Benjamin M. Marx, and Lesley J. Turner. 2019. ProPelled: The effects of grants on graduation, earnings, and welfare. *American Economic Journal: Applied Economics* 11(3): 193-224.

DesJardins, Stephen L., and Brian 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.

Dieterle, Steven, and Andy Snell. 2014. It's hip to be square: Using quadratic first stages to investigate instrument validity and heterogeneous effects. Mimeo.

Dynarski, Susan M. 2002. The behavioral and distributional implications of aid for college. *American Economic Review*, 92(2), 279-285.

Dynarski, Susan M. 2003. 'Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion', *American Economic Review*, 93, 279–88.

Dynarksi, Susan M., & Judith Scott-Clayton. 2013. 'Financial Aid Policy: Lessons From Research', *NBER Working Paper* 18710, National Bureau of Economic Research

Dynarksi, Susan M. 2016. The trouble with student loans? Low earnings, not high debt. Evidence Speaks Report. The Brookings Institute, Washington DC.

Garibaldi, Pietro, Francesco Giavazzi, Andrea Ichino, and Enrico Rettore. 2012 College cost and time to complete a degree: Evidence from tuition discontinuities. *Review of Economics and Statistics* 94(3): 699-711.

Feng, Andy, and Georg Graetz. 2017. A question of degree: the effects of degree class on labor market outcomes. *Economics of Education Review* 61: 140-161.

Goldrick-Rab, Sara, Robert Kelchen, Douglas N. Harris, and James Benson. 2016. Reducing income inequality in educational attainment: Experimental evidence on the impact of financial aid on college completion. *American Journal of Sociology* 121(6): 1762-1817.

Grove, Jack. 2014. 'UK has 'lowest drop-out rate in Europe', Times Higher Education, 1st April 2014. Online. Accessed: 28th July 2017

HEFCE .2013. 'Non-continuation rates at English HEIs', Higher Education Funding Council for England

HESA. 2010. Students and Qualifiers data tables, table 3a, Higher Education Statistics Agency

HESA. 2015. Students and Qualifiers data tables, table 6a, Higher Education Statistics Agency

Joensen, Joanna. 2013. Timing and Incentives: Impacts of Student Aid on Academic Achievement. mimeo, Stockholm School of Economics.

Johnson, Matthew T. 2013. Borrowing constraints, college enrollment, and delayed entry. *Journal of Labor Economics* 31(4): 669-725.

Keane, Michael P., and Kenneth I. Wolpin. 2001. The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review* 42(4): 1051-1103.

Lee, David S. 2009. Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects, *Review of Economic Studies* 76: 1071–1102

Lochner, Lance J., & Alexander Monge-Naranjo. 2012. Credit constraints in education. *Annual Review of Economics*, 4(1), 225-256.

Lochner, Lance J., & Alexander Monge-Naranjo. 2011. The nature of credit constraints and human capital. *American Economic Review*, 101(6), 2487-2529.

Lovenheim, Michael F. 2011. The effect of liquid housing wealth on college enrollment. *Journal of Labor Economics*, 29(4), 741-771.

McCrary, Justin. 2008. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2), 698-714.

Marx, Benjamin M., and Lesley J. Turner. 2018. Borrowing trouble? human capital investment with opt-in costs and implications for the effectiveness of grant aid. *American Economic Journal: Applied Economics* 10(2): 163-201.

Murphy, Richard., Scott-Clayton, Judith., & Gill Wyness. 2019. The end of free college in England: Implications for enrolments, equity, and quality. *Economics of Education Review*, 71, 7-22.

Naylor, Robin, Jeremy Smith, and Shqiponja Telhaj. 2016. Graduate returns, degree class premia and higher education expansion in the UK. *Oxford Economic Papers* 68(2): 525-545.

Nguyen, Tuan D., Jenna W. Kramer, and Brent J. Evans. 2019. The effects of grant aid on student persistence and degree attainment: A systematic review and meta-analysis of the causal evidence. *Review of Educational Research* 89(6): 831-874.

OFFA. 2013. 2011-12 access agreement monitoring: Commitment to widening participation continues', press release, Office for Fair Access

Office of National Statistics. 2012. 'Distribution of median and mean income and tax by age range and gender' *HM Revenue & Customs*, Available at: https://www.gov.uk/government/statistics/distribution-of-median-and-mean-income-and-tax-by-age-range-and-gender-2010-to-2011

Page, Lindsay C., and Judith Scott-Clayton. 2016 Improving college access in the United States: Barriers and policy responses. *Economics of Education Review* 51: 4-22.

Page, Lindsay C., Stacy S. Kehoe, Benjamin L. Castleman, and Gumilang Aryo Sahadewo. More than dollars for scholars the impact of the Dell Scholars Program on college access, persistence, and degree attainment. *Journal of Human Resources* 54, no. 3 (2019): 683-725.

Pollard, Emma, Kari Hadjivassiliou, Sam Swift, and Martha Green. 2017. Credit Transfer in Higher Education: A review of the literature.

Sanderson, Eleanor, and Frank Windmeijer. 2016. A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics* 190(2): 212-221.

Scott-Clayton, Judith., & Susan M Dynarski. 2007. College Grants on a Postcard: A Proposal for Simple and Predictable Student Aid. Hamilton Project Discussion Paper, 1.

Scott-Clayton, Judith. 2011. On money and motivation: A quasi-experimental analysis of financial incentives for college achievement. *Journal of Human Resources*, 46, 614–646.

Scott-Clayton, Judith., & Basit Zafar. 2019. Financial aid, debt management, and socioeconomic outcomes: Post-college effects of merit-based aid. *Journal of Public Economics*, 170, 68-82.

Scrivener, Susan, Michael J. Weiss, Alyssa Ratledge, Timothy Rudd, Colleen Sommo, and Hannah Fresques. 2015. Doubling Graduation Rates: Three-Year Effects of CUNY's Accelerated Study in Associate Programs (ASAP) for Developmental Education Students. New York: MDRC.

Shapiro, Doug, Afet Dundar, Xin Yuan, Autumn T. Harrell, Justin C. Wild, and Mary B. Ziskin. 2014. Some College, No Degree: A National View of Students with Some College Enrollment, but No Completion (Signature Report No. 7). National Student Clearinghouse.

Shea, John. 1997. Instrument relevance in multivariate linear models: A simple measure. *Review of Economics and Statistics* 79: 348–352

SLC. 2013. 'Student Support for Higher Education in England: academic year 2013/14', *Student Loans Company*, November 28th 2013.

Stinebrickner, Ralph., & Todd Stinebrickner. 2008. The effect of credit constraints on the college drop-out decision: A direct approach using a new panel study. *American Economic Review*, 98(5), 2163-84.

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

UCAS .2022. 'Replying to your UCAS undergraduate offers', University College Admissions System, available at: https://www.ucas.com/undergraduate/after-you-apply/types-offer/replying-your-ucas-undergraduate-offers#changed-your-mind

Van der Klaauw, Wilbert. 2002. Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression-Discontinuity Approach. *International Economic Review*, 43(4), 1249-1287.

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

Wooldridge Jeffrey M. 2009. *Econometric Analysis of Cross Section and Panel Data*. The MIT Press Cambridge, Massachusetts.

Figure 1: Financial Aid Schedules at Universities over Time

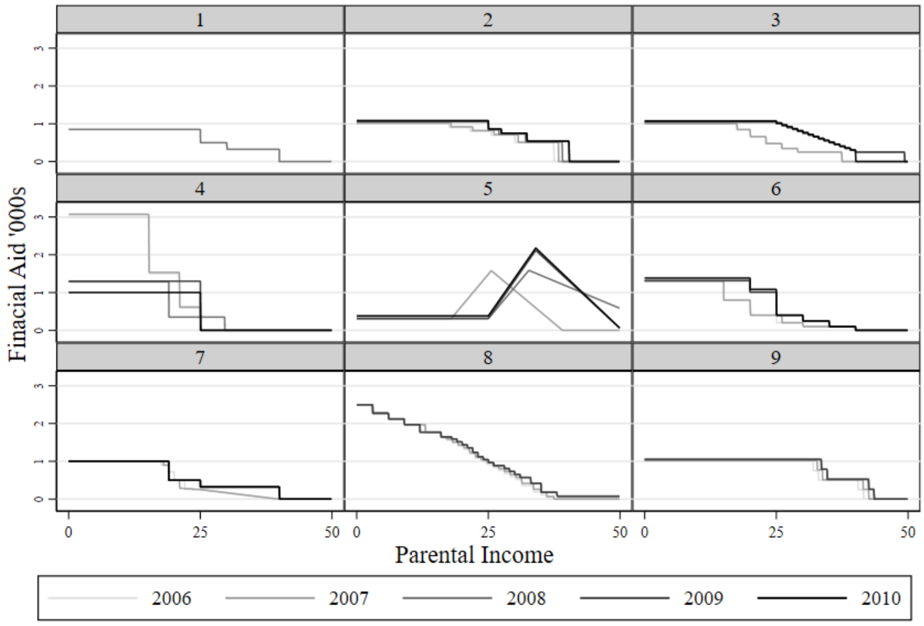

Notes: This presents the Institutional Financial Aid (IFA) schedules for first year students for nine anonymous universities for students entering in the years 2006 through to 2010. Figures reported in nominal values. Parental income in thousands of pounds.

Figure 2: IFA Rules and Compliance at Universities in 2008

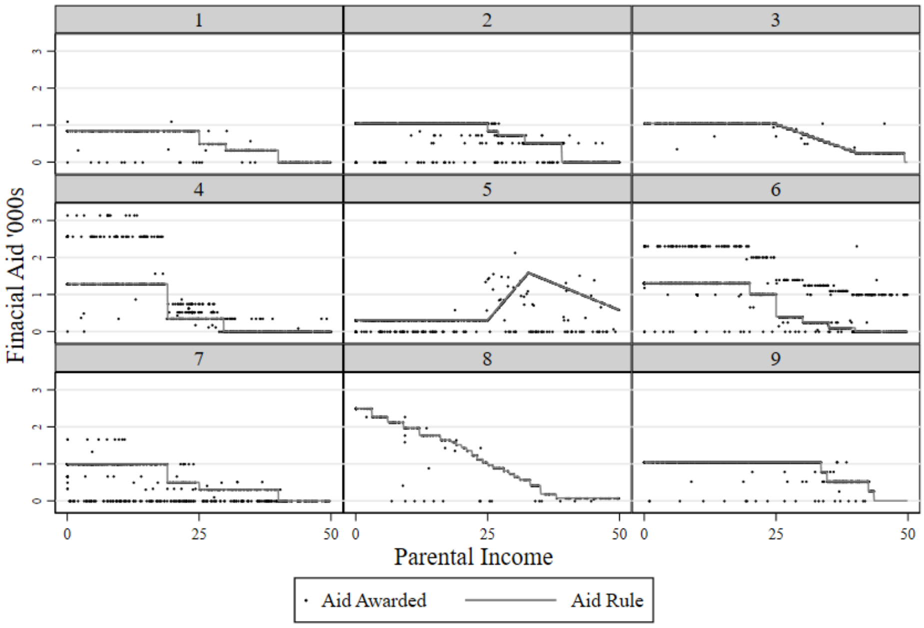

Notes: This figure shows household income and IFA receipt for every first-year student in 2008 for each university. University 1 shows the compliance in 2010 rather than 2008, as that year of entry is not available for that university. The horizontal and vertical lines show the different bursary levels advertised by the university at each income level. Parental income in thousands of pounds. Each point represents an individual's eligible bursary amount.

Figure 3: First Year Test Scores Residuals, when excluding IFA, by Distance from Cut-Off

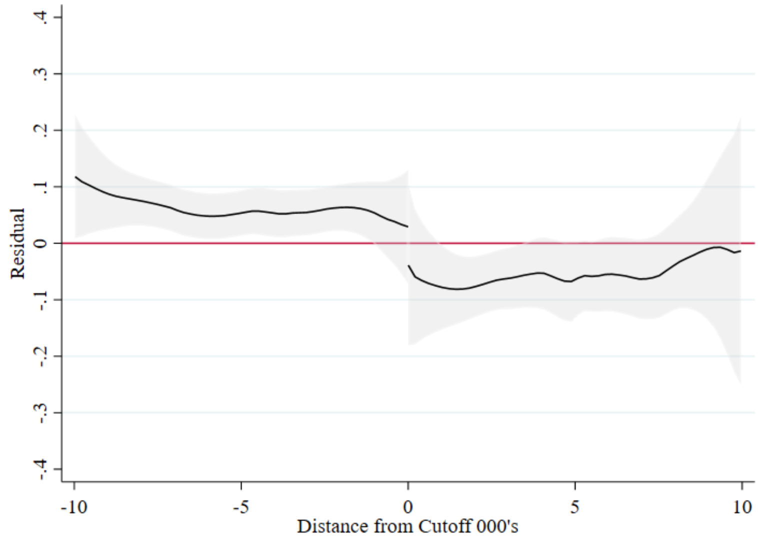

Notes: This figure plots the residuals from the main specification (2) excluding the IFA parameters, plotted according to distance from cut-off. This is presented for parental incomes up to £10,000 above and £10,000 below the nearest cut-off) for discrete changes in bursary amounts over £100 (n=9093). Plotted using the lpolci command using a local linear regression estimator, using Epanechnikov kernel weighting with shaded area representing the 95 percent confidence interval.

Figure 4: First Year Test Scores Residuals, when excluding IFA, by Financial Aid Amount

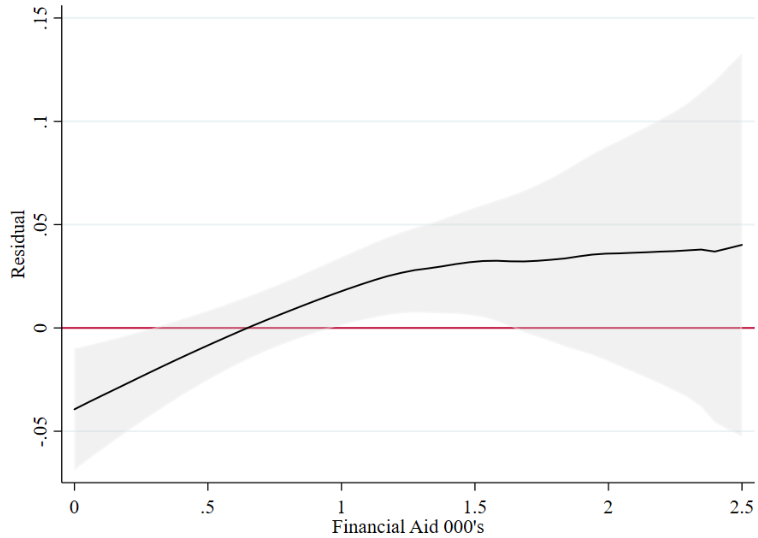

Notes: This figure shows the residuals from the main specification (2) excluding the IFA parameters against IFA received using the main sample (n=17,060). Plotted using the lployci command, using a local linear regression estimator, with a 95 percent confidence interval.

Figure 5: First Year Test Scores Outcome Maximisation and Cost Minimisation of IFA

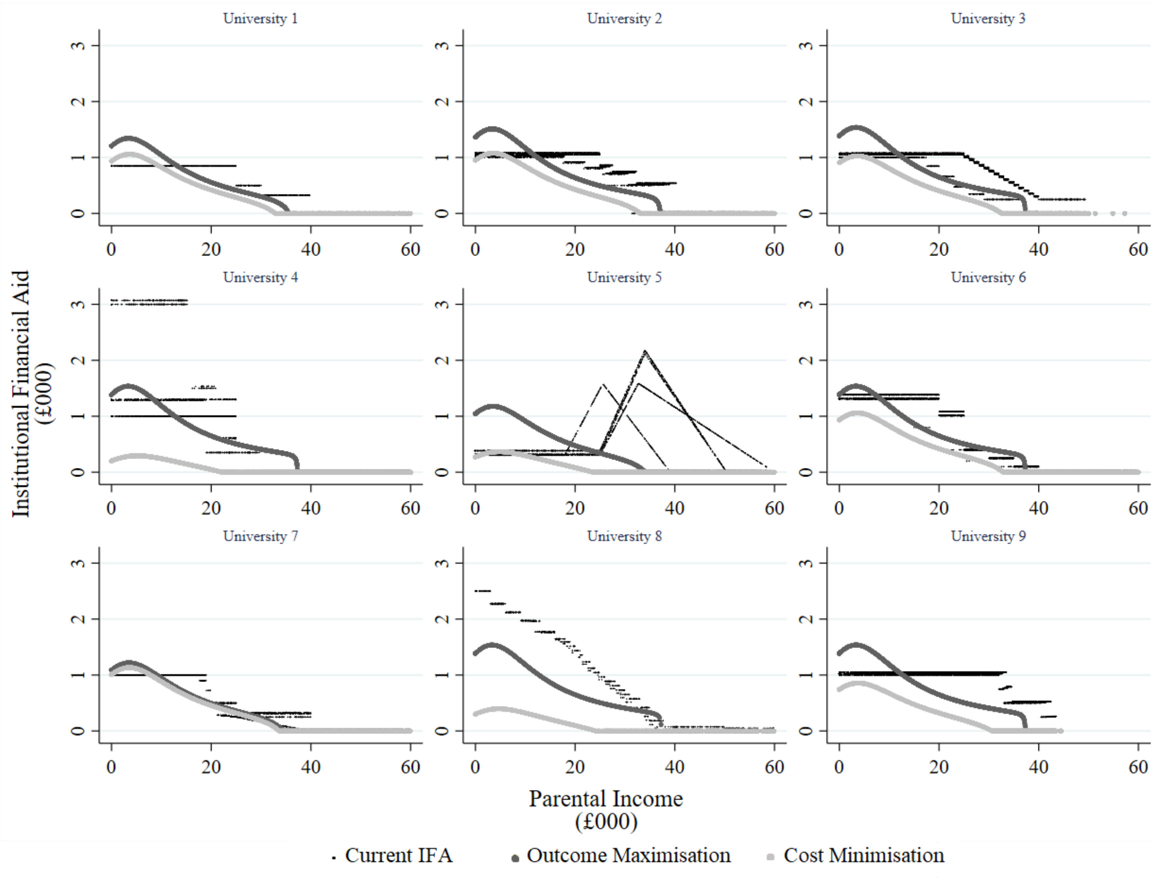

Notes: This figure presents the results of an optimisation for each university given their student population and our estimates of the impact of aid. The black points represent the eligible IFA for current students using IFA rules applicable to them. The dark grey circles represent IFA policies maximising first year test scores holding spending constant. The light grey circles represent IFA policies minimising spending, holding first year completion rates constant. Estimates are from a specification of a quadratic in aid and a cubic of parental income and their interactions, conditional on university-year effects, student demographics. See Online Appendix E for details.

Table 1: Student Characteristics and Outcomes

	All Students		Balanced Panel of Students	
	(1)	(2)	(3)	(4)
	Mean	Std Dev	Mean	Std.dev
Household Income	£23,261	£19,476	£23,288	£19,253
Maintenance Grant	£2,059	£1,076	£2,035	£1,055
Eligible IFA (E)	£759	£595	£787	£637
IFA Received (A)	£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				
First Year	0.95	0.22	0.95	0.21
Second Year	0.92	0.27	0.92	0.27
Third Year	0.89	0.31	0.9	0.3
Standardised Scores				
First Year	0.00	1.00	0.03	0.96
Second Year	0.00	1.00	0.03	0.96
Third Year	0.00	1.00	0.04	0.95
Obtain Good-Degree	0.60	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. Standardized test score data only available for 17,060 students with a Balanced Panel. Data on students' household income is observed for all nine universities shown.

Table 2: Impact of First Year IFA on Student Performance

	Complete First Year (1)	Course Scores First Year (2)	Obtain Good Degree (3)
Panel A - OLS			
IFA (A_i)	0.035	0.039	0.036
	(0.014)	(0.020)	(0.012)
Panel B - OLS			
IFA (A_i)	0.140	0.183	0.131
	(0.035)	(0.047)	(0.025)
IFA squared (A_i^2)	-0.042	-0.057	-0.038
	(0.011)	(0.014)	(0.009)
IFA Marginal Effect (A_i)	0.074	0.091	0.072
	(0.019)	(0.024)	(0.013)
R-squared	0.136	0.066	0.142
Panel C – Reduced Form			
Eligible IFA Marginal effects (E_i)	0.012	0.050	0.029
	(0.006)	(0.017)	(0.011)
Panel D – 2SLS			
IFA $(\widehat{A_i})$	0.029	0.130	0.056
	(0.011)	(0.032)	(0.022)
IFA Squared $(\widehat{A_i^2})$	-0.010	-0.044	-0.014
	(0.003)	(0.010)	(0.007)
Marginal Effect $(\widehat{A_i})$	0.014	0.059	0.034
	(0.006)	(0.017)	(0.012)
R-Squared	0.125	0.066	0.141
First Stage $F_{Aid\|Aid\text{^}2}$	1340	676	1340
First Stage $F_{Aid\text{^}2\|Aid}$	2203	751	2203
Sheas's Adj-P R^2 Aid	0.458	0.526	0.458
Sheas's Adj-P R^2 Aid^2	0.531	0.570	0.531
Observations	22,770	17,060	22,770
University-Year Effects θ_{jt}	✓	✓	✓
Student Characteristics	✓	✓	✓
$f_t(I_i)$	✓	✓	✓

Notes: IFA is the amount of Institutional Financial Aid received during the first year measured in £1,000's. Assumes parental income takes the function of a fourth order polynomial, which can vary by year of entry, $f_t(I_i)$. 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. Coefficients in panel C are average marginal effect. Sample consists only of those students whose final outcome can be observed. Standard errors are in parentheses, and are clustered at institution-year level. First Stage F-statistics are Conditional F-Statistics accounting for the multiple endogenous variables (Sanderson and Windmeijer, 2014). Sheas's Adj-P represents Shea's adjusted partial R2 statistic (Shea, 1997).

Table 3: Average Marginal Impact of Running Average of IFA

Marginal effects	Complete First Year (1)	Complete Second Year (2)	Complete Third Year (3)	Course Scores First Year (4)	Course Scores Second Year (5)	Course Scores Third Year (6)	Obtain Good Degree (7)
Panel A – Reduced Form							
Running Average Eligible IFA (E_i)	0.012 (0.006)	0.015 (0.006)	0.016 (0.007)	0.050 (0.017)	0.031 (0.014)	0.014 (0.023)	0.029 (0.011)
Panel B – 2SLS							
Running Average Received IFA $(\widehat{A_i})$	0.014 (0.006)	0.021 (0.008)	0.027 (0.010)	0.059 (0.017)	0.041 (0.018)	0.019 (0.032)	0.045 (0.018)
Observations	22,770	22,770	22,770	17,060	16,218	14,623	22,770
University-Year Effects θ_{jt}	✓	✓	✓	✓	✓	✓	✓
Student Characteristics	✓	✓	✓	✓	✓	✓	✓
$f_t(I_i)$	✓	✓	✓	✓	✓	✓	✓

Notes: This table presents the average marginal effect form 14 specifications. Running Average of Eligible IFA amount is the average amount of bursary aid student would be eligible for over the corresponding number of years regardless of enrolment status, based on first year parental income and the concurrent eligibility rules. Assumes parental income takes the function of a fourth order polynomial, which can vary by year of entry, $f_t(I_i)$. Test score information is not available for all university years. Sample consists only of those students whose final outcome can be observed. Standard errors are in parentheses, and are clustered at institution-year level.

Table 4: Alternative specifications – Average Maringal Effects on First Year Course Scores

Cohort indicators with parental income polynomials	University indicators with parental income polynomials				
	$\mathbf{1}_j * I_i^0$	$+\mathbf{1}_j * I_i^1$	$+\mathbf{1}_j * I_i^2$	$+\mathbf{1}_j * I_i^3$	$+\mathbf{1}_j * I_i^4$
	(1)	(2)	(3)	(4)	(5)
$\mathbf{1}_t * I_i^0$	0.068	0.052	0.071	0.057	0.012
	(0.015)	(0.022)	(0.026)	(0.023)	(0.035)
	-	*0.704*	*0.119*	*0.167*	*0.873*
$+\mathbf{1}_t * I_i^1$	0.072	0.051	0.071	0.057	0.011
	(0.016)	(0.022)	(0.025)	(0.023)	(0.035)
	-	*0.774*	*0.133*	*0.117*	*0.882*
$+\mathbf{1}_t * I_i^2$	0.072	0.062	0.070	0.060	0.015
	(0.015)	(0.016)	(0.025)	(0.022)	(0.034)
	-	*0.621*	*0.136*	*0.117*	*0.882*
$+\mathbf{1}_t * I_i^3$	0.067	0.056	0.068	0.056	0.014
	(0.015)	(0.016)	(0.026)	(0.025)	(0.036)
	-	*0753*	*0.190*	*0.326*	*0.870*
$+\mathbf{1}_t * I_i^4$	0.059	0.045	0.050	0.039	0.017
	(0.017)	(0.018)	(0.025)	(0.027)	(0.037)
	-	*0.435*	*0.188*	*0.315*	*0.886*

	Parental income polynomials				
	$\mathbf{1}_{jt} * I_i^0$	$+\mathbf{1}_{jt} * I_i^1$	$+\mathbf{1}_{jt} * I_i^2$	$+\mathbf{1}_{tj} * I_i^3$	$+\mathbf{1}_{tj} * I_i^4$
	(1)	(2)	(3)	(4)	(5)
University-Cohort indicators	0.068	0.053	0.073	0.052	0.005
	(0.015)	(0.022)	(0.026)	(0.026)	(0.039)
	-	*0.875*	*0.515*	*0.408*	*0.956*

Notes: Values presented are average marginal effects, using the quadratic specification for aid. Where $\mathbf{1}_t$ is an indicator function for entry cohort year t, $\mathbf{1}_j$ is an indicator function for university j, and $\mathbf{1}_{jt}$ is an indicator function for university-entry cohort jt. All specifications additionally include university entry-cohort fixed effects, θ_{jt}, and student characteristics. Standard errors are in parentheses, and are clustered at institution-year level. P-values from Chi2 distribution in are presented in italics testing the joint significance of the additional parameters relative to the preceding column.

Table 5: Heterogeneity of Average Maringal Effects

Outcome	All	Low Entry Test Scores	High Entry Test Scores	25th Percentile of Parental income	50th Percentile of Parental income	75th Percentile of Parental income
	(1)	(2)	(3)	(4)	(5)	(6)
Complete First Year	0.014	0.011	0.034	0.020	0.016	0.014
	(0.006)	(0.005)	(0.008)	(0.010)	(0.006)	(0.009)
Diff. P-Value		0.015		0.098		0.098
Mean	[0.95]	[0.95]	[0.98]	[0.948]	[0.951]	[0.958]
Course Scores First Year	0.059	0.038	0.072	0.083	0.063	0.038
	(0.017)	(0.021)	(0.038)	(0.025)	(0.020)	(0.024)
Diff. P-Value		0.433		0.130		0.130
Mean	[0.03]	[-0.07]	[0.18]	[-0.023]	[0.045]	[0.076]
Good-degree	0.034	0.026	0.085	0.039	0.033	0.028
	(0.012)	(0.014)	(0.015)	(0.015)	(0.010)	(0.013)
Diff. P-Value		0.004		0.555		0.555
Mean	[0.63]	[0.56]	[0.67]	[0.59]	[0.64]	[0.68]
N	22,770	9,795	7,733	22,770	22,770	22,770

Notes: All specifications are conditional on fourth order polynomial, which can vary by year of entry $f_t(I_i)$, university-year effects θ_{jt}, and student characteristics. Coefficients presented are using the quadratic specification for aid, and are average marginal effects. Sample consists only of those students whose final outcome can be observed. Columns 2 and 3 are estimated using mutually exclusive sub-samples as defined by column titles. Low Tariff is defined by any student under the 50th entry test score percentile. For these columns Diff. P-Value represents the probability that the coefficients of the subsamples are significantly different. Columns 4 through 6 use the whole sample, the parameters are the estimated marginal effects evaluated at different levels of parental income. For 25th, 50th and 75th parental income percentiles, mean value taken for 0-50th, 25-75th, and 50-100th percentiles respectively. For these columns Diff. P-Value represents the probability that the marginal effects at the 25th and 50th percentile are different and the marginal effects at the 50th and 75th percentile are different. The 5,242 students with no recorded entry test scores are excluded from the test score heterogeneity. Standard errors are in parentheses, and are clustered at institution year level. Mean value of outcome in square parentheses.

Table 6: Optimized versus Current IFA Spending – University Average First Year Course Scores

University	1	2	3	4	5	6	7	8	9
First Year Test Scores	-0.001	0.004	0.031	0.010	0.038	0.096	0.000	0.015	0.023
Gains from Maximising Outcomes s.t. Current Spending	0.005	0.007	0.008	0.030	0.034	0.006	0.002	0.037	0.014
% Savings from Cost Minimization s.t. Current Outcomes	25.3	32.9	43.8	89.2	74.1	43.2	7.6	88.0	61.2
Cost Minimised Savings per student (£,000)	0.135	0.255	0.378	0.627	0.406	0.326	0.055	1.166	0.561

Notes: Derived values obtained from specification described in Appendix E: Optimisation, and from outcome maximisation E2 and cost minimisation E3. The first-row are the current standardized test scores, by university prior to sample restrictions. The second-row are the gains from changing aid with regards to maximising test scores. The third-row is the proportion of current spending that would be saved if spending was minimised, while maintaining test scores. The fourth-row are the savings per student in £,000s from adopting a cost minimising scheme while keeping aggregate outcomes constant.