NBER WORKING PAPER SERIES

DO INCOME CONTINGENT STUDENT LOAN PROGRAMS DISTORT EARNINGS? EVIDENCE FROM THE UK

Jack W. Britton
Jonathan Gruber

Working Paper 25822
http://www.nber.org/papers/w25822

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2019

Corresponding author, Jack Britton (jack.britton@ifs.org.uk). Britton would like to thank the British Academy for funding this work, grant PF150081 and the Economic and Social Research Council (ESRC)'s Centre for the Microeconomic Analysis of Public Policy at IFS (grant reference ES/M010147.1) for financial support. This work contains statistical data from Her Majesty's Revenue and Customs (HMRC) which are Crown Copyright. The research data sets used may not exactly reproduce HMRC aggregates. The use of HMRC statistical data in this work does not imply the endorsement of HMRC in relation to the interpretation or analysis of the information. The data creators, depositors, copyright holders and funders bear no responsibility for the analysis, inferences, conclusions or interpretation of the data presented here. Responsibility for interpretation of the data, as well as for any errors, is the authors' alone. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Do Income Contingent Student Loan Programs Distort Earnings? Evidence from the UK
Jack W. Britton and Jonathan Gruber
NBER Working Paper No. 25822
May 2019
JEL No. H2,H52,I22

ABSTRACT

Government backed income contingent student loans are an increasingly being used to fund higher education. An income contingent repayment plan acts as an incremental marginal tax on labor earnings, which could cause individuals to distort their work effort. This paper uses an administrative dataset from the UK that links student loan borrowers between 1998 and 2008, to their official tax records between 2001/02 and 2013/14. Using a combination of techniques, including bunching and difference-in-difference methodology, our findings strongly reject the hypothesis that the UK's income-contingent repayment plan distorts labor supply.

Jack W. Britton
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
United Kingdom
jack_b@ifs.org.uk

Jonathan Gruber
Department of Economics, E52-434
MIT
50 Memorial Drive
Cambridge, MA 02142
and NBER
gruberj@mit.edu

The rising cost of higher education is a topic of enormous controversy in the U.S., and increasingly in other nations as well. In the U.S., the average cost of a four year private college degree has risen from 15,160 in 1987 to $35,740 today.[2] In other countries, in contrast, higher education is largely free historically, but in a number of countries the costs of higher education are rising as well.[3]

A primary source of financing these higher education expenses is government subsidized student loans. In the U.S., qualifying undergraduate students can borrow up to $31,000 ($23,000 in subsidized loans) per degree program. Thirteen million students use roughly $120 billion per year in student aid.[4]

The problem with this solution is that students often appear unable to make their loan payments upon graduation. For college graduates in 2007-2008 who had not pursued further education, 24% were delinquent on at least one payment, 5% had defaulted on their debts, and 22% had debt repayments which amounted to more than 12% of their income, a level considered by experts to be burdensome.[5] The total burden of student debt in the U.S. is over $1.4 trillion.[6] This has led to a notable public debate about the college "affordability crisis" in the U.S., with proposed solutions ranging from limiting the set of institutions at which students can obtain loans to making universities debt free for all students.[7]

---

[2] https://trends.collegeboard.org/sites/default/files/2017-trends-in-college-pricing_0.pdf
[3] https://read.oecd-ilibrary.org/education/education-at-a-glance-2018_eag-2018-en#page304
[4] https://studentaid.ed.gov/sa/about
[5] https://nces.ed.gov/pubs2017/2017436.pdf

[6]https://www.vox.com/policy-and-politics/2018/3/20/17124080/brian-schatz-debt-free-college-plan-senate

[7]The Obama administration proposed denying student loan eligibility to some colleges whose graduates did not earn enough to pay off their student debt; the Trump administration has delayed implementation of this rule. https://www.nytimes.com/2017/06/14/business/student-loans-for-profit-schools-colleges.html and https://www.nytimes.com/2018/05/13/business/education-department-for-profit-colleges.html. At the same

But another option is to emphasize a solution that is already in place: more use of income contingent loans. Currently, undergraduate students can borrow up to around $60,000 ($140,000 for graduate students) in federal student loans and by default enter a "mortgage style" repayment system, which means repayments are set at a fixed level each year, with the level chosen so that the debt is cleared in 10 years. However, since 2009, borrowers have been able to apply to switch to Income Based Repayment (IBR), which have a maximum contribution of 15% of discretionary income and 25 year repayment period, after which any outstanding debt is written off.[8] In 2016, 28% of borrowers were enrolled in income driven loan repayment programs, holding 48% of all student loan debt.[9]

An income-contingent loan approach to financing college education poses the standard tradeoff of any means-tested transfer: insurance against poor labor market outcomes versus disincentives to earn. An income-contingent repayment plan acts as an incremental marginal tax on labor earnings, which could cause individuals to distort their work effort. This could be particularly true for workers in the lower-middle class, who already face fairly high tax rates from the phaseout of other government entitlement programs. Whether income-contingent loan programs have such a distortionary effect on labor supply therefore becomes a critical question for evaluating the advisability of this financing approach.

---

time, political leaders have proposed debt free college:https://www.vox.com/policy-and-politics/2018/3/20/17124080/brian-schatz-debt-free-college-plan-senate

[8] See Mueller and Yanellis (2018). People can switch in IBR if their income to debt ratio is sufficiently low.

[9] https://trends.collegeboard.org/sites/default/files/2017-trends-student-aid_0.pdf

Assessing the impact of income-contingent loans on labor supply decisions is difficult in the U.S. Income-contingent loans are an option, and the choice of that option is likely correlated with earnings prospects and labor supply decisions.

So to study the effect of contingent loan programs on labor supply we turn to England. College in England was historically very low cost, but this has increased significantly with time, with fees rising from around £1,000 per year to around £3,000 per year in 2006, to £9,000 in 2012; as well, there have been sizeable increases in the cost of living, particularly for those coming from poorer household or attending university in London. The government has issued student loans to help with college costs for many years. Since 1998, the share of eligible individuals choosing to borrow has remained relatively stable at around 85%. Private loans during this period were not at all common due to the income contingent nature of the loans and the low interest rate offered – until 2012, interest was equal to the minimum of the Retail Price Index and the Bank Rate + 1%.

Most importantly, repayment of income-contingent loans is only made once earnings has crossed a certain threshold. This implies a large marginal tax rate at that threshold that can be studied using standard bunching techniques (Saez, 2010). Moreover, this earnings threshold changed repeatedly during our sample, allowing us to rule out other reasons for possible bunching at the threshold. And the change in threshold allows us to also consider an alternative differences-in-differences empirical strategy that compares earnings changes in income ranges that see differential changes in income and substitution effects as the tax threshold changes. This allows us to confirm our findings using multiple approaches.

We do so using a dataset that links Student Loan Company (SLC) records to Her Majesty's Revenue and Customs (HMRC) data. This consists of people who lived in England when they applied for university and borrowed from the SLC to attend university in the UK. We have a 10% sample of everyone who borrowed from the SLC between 1998 and 2008, and all of their tax records between 2001/02 and 2013/14. We also have access to a residual dataset that consists of all of the tax records that do not match to the SLC data and consists mostly of people who didn't go to university, or those who did go but did not borrow, which allows us to form a control group for our analysis.

Our findings are clear: there is no meaningful impact on earnings from the income-contingent repayment rules. Across a variety of thresholds, we see little consistent evidence of bunching. And we do not find any consistent change in earnings in income ranges that are impacted through the change in the taxable threshold. This result is robust across types of individuals and types of income, both payroll-based and self-reported. Our findings suggest that there is little labor supply distortion to using this mechanism to finance college education.

Our paper proceeds as follows. Part I discusses the issues around income contingent loans as a college financing tool. Part II reviews the system of college financing in the U.K. Part III discusses our data and Part IV our empirical strategy. Part V presents results, and Part VI concludes.

Part I: Income-Contingent Loans as a College Financing Tool

A large literature in economics documents the high rate of return to college attendance, as reviewed by Hout (2012) and Oreopoulos and Petronijevic (2013). At the same time, this is

only an average estimate, and this human capital investment has uncertain returns. Risk averse consumers facing such an enormous investment with uncertainty will value insurance against those uncertain returns. For this reason, income contingent financing has long been recommended as an approach, dating at least from Friedman and Kuznets (1945); see Shireman (2017) for an overview of the history of thought in this area.

But as with any insurance scheme, the benefits of reduced uncertainty must be weighed against potential moral hazard costs. In this case, the moral hazard is reduced income generation in response to the higher implicit taxes put in place by an income-contingent system.

These considerations may be paramount in countries such as the U.S. or the U.K. which rely heavily on income-contingent transfers for the lower part of the income distribution. Programs such as the Earned Income Tax Credit in the U.S. or the Working Tax Credit in the U.K. have proven to be effective tools to redistribute income while promoting entry into the labor force. At the same time, marginal tax rates in the phase out range of these programs can be large. This is further augmented by other income-based transfers such as housing, health insurance, and food assistance.[10]

The available evidence on income responses to marginal tax rates under existing redistribution programs does not suggest that such distortions are in fact sizeable.[11] The relevance of this literature in this context is not clear, however. The populations eligible for

[10] For an example of the analysis of these types of high marginal rates, see Mulligan (2015).
[11] See for example limited evidence on the distortionary role of marginal tax rates under the EITC in Eissa and Liebman (1996) or more recent work on the limited impact of high marginal rates on full time work under the ACA in Moriya, Selden, and Simon (2016).

these welfare programs are largely not college attendees. And many of the implicit taxes under these programs are difficult to compute and do not tie to an easily identifiable level of income.

Perhaps more relevant is the literature on income responses to tax thresholds. The majority of which has focused on those with high incomes. Gruber and Saez (2002) estimate an elasticity of taxable income (ETI) of 0.57 for high income individuals (those with incomes above $100,000 in the 1980s) in the US. Brewer, Saez, and Shephard (2010) exploit a large cut in the top rate of tax in the UK from 83% to 40% to estimate an ETI of 0.46 for those in the top 1% of the income distribution.[12]

Relatively few papers look at the responsiveness further down the income distribution. Gruber and Saez (2002) argue that the effects are much smaller for those with incomes below $100,000 in the US, while Adam et al. (2017) in the UK and Saez (2010) in the US find that employees – who make up the vast majority of their sample – do not bunch at kinks in the tax schedule further down the income distribution.

Chetty (2012) highlights the differences between macro and micro elasticities of labor supply and argues that this could largely be driven by optimizing frictions. Gruber and Saez (2002) find that taxpayers who itemize (and who are likely to have the lowest frictions) are particularly responsive to taxation, Kleven and Waseem (2013) find that the frictions are large in Pakistan, while Adam et al. (2017) estimate that they could be worth as much as 9% of income based on evidence from the UK. They also show that the most responsive to taxes are company owner-managers and the self-employed. We are able to estimate responses of those

[12] These estimates seem high compared to the much of the more recent literature. Saez, Slemrod, and Giertz (2012) suggest that the midpoint of the range of best available estimates is around 0.25. Piketty, Saez, and Stantcheva (2014) also estimates of a very similar magnitude to this.

who itemize (that is, those who file Self-Assessment tax forms) separately from those who do not, and we find no evidence of effects for either group.

Once again, the applicability of this tax literature to modeling the effect of income-contingent loan payoffs isn't clear. That literature has not focused explicitly on the recent college graduates the relevant income range for considering labor supply response to repayment thresholds. College graduates who are undergoing frequent job transitions may be more able to adjust their incomes. And these repayment thresholds may be even more salient than are kinks in the income tax schedule.

The literature that explicitly looks at the labor supply responses to student loans is extremely limited. [13] The exception is Chapman and Leigh (2009), who investigate bunching at the student loan thresholds in Australia in the early 2000s. At this time, Australian students who borrowed from the Higher Education Contribution Scheme (HECS) had to make repayments of 4% of their taxable income once they earned above the threshold of around AUS $24,000. Crucially, this is a (vary large) kink in schedule which meant those earning exactly the threshold would have disposable income around AUS $760 lower than someone earning AUS $1 less. They argue that there is evidence of bunching at the threshold, suggesting graduates are indeed able to adjust their incomes in response to thresholds. However, they also argue that while the effect is statistically significant, it is economically small – they suggest only around 0.3% of all those with HECS debt bunch, and they predict that the degree to which they reduce

[13] Also relevant is the literature investigating the impact of student loans on career choices. Rothstein and Rouse show that people shift into higher paid occupations, while Sieg and Wang (2018) show that student loans affect marriage and career choices of lawyers. The literature on the impact of income contingent loans is still in its infancy, however.

their income is likely to be small. Our paper importantly extends on this work in several ways. First, we use a much large sample (a 10% cut of administrative data) that gives us many more individuals near to the threshold; Chapman and Leigh (2009) have only around 600 people earning within AUS $1000 of the threshold in total, while we have several thousand individuals within that range of the threshold *in each year*, enabling us to look separately at multiple thresholds over multiple years. Second, they use a simple probit model to investigate the degree of excess mass caused by the notch in the schedule, but they do not calculate elasticities. And third, they are not able to consider how earnings growth in the early part of graduates' careers is affected by the thresholds, which we are able to do by exploiting the changes in the threshold and through the panel element of our dataset.

Part II: College Financing in the U.K.

The higher education system in the United Kingdom is complicated and has been subject to much reform over the past 20 years. The U.K. is divided into four devolved nations - England, Wales, Scotland and Northern Ireland - and higher education policy is the responsibility of these devolved governments. This means that policy varies considerably. The focus of this paper is on England, which is the country for which we have obtained Student Loan Company (SLC) records of borrowers.

English students can choose between around 150 universities and several other "further education institutions" anywhere in the UK for which they are eligible for student loans. The vast majority of these institutions remain as public institutions; although more recently the number of private institutions has grown, this recent growth is not relevant to this paper as it

occurred after the population of students that we study had attended university. Traditionally, the public institutions charged very low tuition: until 1998, there were no fees at all; between 1998 and 2005 fees were around £1,000; between 2006 and 2011 they were around £3,000; then in 2012 they increased to £9,000. Details are provided in Table 1.

Table 1: English student loan amounts and repayment thresholds by tax year

Year	Maintenance Loan	Tuition fees	Repayment threshold
1998/99	£2,735	£1,075	n/a
1999/00	£3,635	£1,075	n/a
2000/01	£3,725	£1,075	£10,000
2001/02**	£3,815	£1,075	£10,000
2002/03	£3,905	£1,100	£10,000
2003/04	£4,000	£1,125	£10,000
2004/05	£4,095	£1,150	£10,000
2005/06	£4,195	£3,000*	£15,000
2006/07	£4,405	£3,000*	£15,000
2007/08	£4,510	£3,070*	£15,000
2008/09	£4,625	£3,145*	£15,000
2009/10	£4,950	£3,225*	£15,000
2010/11	£4,950	£3,290*	£15,000
2011/12	£4,950	£3,375*	£15,000
2012/13	£5,500	£9,000*	£15,795
2013/14	£5,500	£9,000*	£16,365

Notes: Cash terms figures. The maintenance loan amounts show the maximum maintenance loan for somebody living away from home, outside of London. London rates are higher, while living at home rates are lower (generally, those from richer households can borrow less). The repayment threshold is for the tax year, which runs from April 6-April 5. Maintenance and tuition fee loans are for September 1 – August 31. The repayment threshold apply to people in repayment, meaning they must have left university. Thus, people starting university in 2003/04 and doing a 3 year course would enter repayment in 2006/07, meaning they would face the £15,000 repayment threshold. The repayment rate is 9% of earnings above the threshold throughout. Source: Student Loan Company Statistical Release Series "Student support for higher education in England". *Indicates tuition fee loan availability **Indicates the first year for which we have earnings data.

The government has offered student loans to help with living expenses during study since the early 1980s. It set up the "Student Loan Company" in 1990 to administer the loans, and in 1998 it made all newly issued loans "income contingent" which meant that individuals only had to repay them if they earned above a certain threshold (prior to this, loans were "mortgage-style" meaning repayments were dependent on amount borrowed rather than

current income). These living cost loans (referred to as "maintenance" loans) were around £4,000 in 2002/03, increasing roughly in line with inflation. In 2006, for the first time the government issued additional tuition fee loans, which almost doubled the amount people were able to borrow.[14]

The income contingent loan that was first introduced for 1998 university entrants meant that individuals had to repay 9% of all income they earn above a threshold set by the government. This was initially set at £10,000, before being increased to £15,000 in 2005, and then roughly indexed to inflation after 2012. The marginal tax rate at the point just below this repayment threshold was 33% between 2001/02 and 2008/09 and 31% between 2009/10 and 2013/14, the last year we observe the income data. This consisted 11% in National Insurance Contributions (NICs), with the remainder made up in income tax. This meant that student loans contributions increased the marginal tax rate from 33% to 42% up until 2008/09 and from 31% to 40% after 2008/09. The student loan thresholds were chosen to be above the range where they would interact with the phase out of other social support programs.

There have also been changes to the details of the loan package such as the repayment period and interest rate. Loans for our earlier cohorts (1998-2005) were written off at age 65, while for our later cohorts (2006-2008) they were written off 25 years after graduation. Meanwhile, the interest rate on student debt was index linked to be equal to the minimum of RPI and the Bank of England base rate plus 1%, meaning it was much higher before the major recession in 2008. There have been further subsequent changes to the repayment period

[14] Before 2006, the government did not offer explicit tuition fee loans, which meant students had to fund their own tuition privately, or alternatively use their maintenance loans on their tuition costs. The tuition fees during the 1998-2005 period were means-tested which meant that poorer students were exempt from paying any fees.

(increased to 30 years) and the interest rate (increased to RPI plus 3% while studying and RPI plus 0-3% depending on income after graduation), but these occurred for cohorts starting university from 2012, who we do not observe in our sample.

Despite the changes to the loans over the past 20 years, the share of individuals choosing to borrow from the SLC has remained remarkably stable at around 85% (Britton, Shephard and Vignoles, 2018). Amongst non-borrowers, most have their university education funded by their parents, and while private borrowing for student finance in the UK is very rare.[15]

Student loan repayments are automatically deducted by the tax authority, HMRC, from individual's reported income. Amongst our sample of borrowers who are in contact with HMRC, around 90% have their income reported to HMRC through Pay As You Earn (PAYE) tax forms, which are automatically submitted by their employers. The other 10% file Self-Assessment (SA) tax forms which include income from employment, self-employment and partnerships as well as other forms of income (see Britton, Shephard and Vignoles, 2018 for more detail). These individuals are legally required to accurately report their income in these forms. The loan repayments are taken based on gross income net of salary sacrifice pension contribution in the UK tax year, which runs from April 6th to April 5th. HMRC informs the SLC about the repayments made by each individual at the end of the tax year (SLC typically contacts individuals when they are likely to clear before the end of the tax year so the borrower can send in evidence of repayment before the end of the tax year to prevent overpaying). Individuals are eligible to start making repayments on their student loan only in the first full tax year after they have left

[15] Author's own calculations using the Longitudinal Survey of Young People in England.

college. In most cases, students graduate in the summer and then start making repayments based on the earnings from the following April.

Part III: Data

Our data for this analysis comes from Student Loan Company (SLC) records linked to Her Majesty's Revenue and Customs (HMRC) tax records (see Britton, Shephard and Vignoles, 2018 for more detail). We have access to a 10% random sample of all students who borrowed from the English SLC between 1998 and 2008. This corresponds to around 24,000 individuals, around 55% of whom are female, reflecting the higher share of women who attend higher education in England. Anyone who was domiciled in England upon application to higher education (typically in the year before they start university) and attended an institution anywhere in the UK is eligible to borrow, and we estimate that around 85% of those eligible do choose to borrow. The SLC and HMRC databases are hard linked based on official National Insurance numbers (equivalent to Social Security numbers in the US), which are de-identified to ensure researchers cannot observe them.

From the SLC data we observe total amount borrowed, "voluntary repayments" and from the HMRC data we observe earnings. Voluntary repayments are payments made over and above those required based on income. For example, in a year where the threshold is £10,000, someone with an income of £15,000 would be obligated to pay 0.09*(£15,000-£10,000) = £450. However, they can make voluntary repayments over and above that figure which we observe (more detail on these is documented in Britton et al., 2018b). Combining these factors, we are able to estimate outstanding student debt for each individual at the end of a given tax year in our linked dataset. That is, we calculate compulsory repayments based on observed earnings,

take off voluntary repayments and add loan interest. Figure 1 shows our estimated number of borrowers who are in the labor market and have outstanding debt by tax year, comparing our estimates from the HMRC dataset with official records from the SLC. Since we are using a 10% sample, we divide the SLC numbers by ten for direct comparability. The downward tick towards the end occurs because there are no new graduates entering the labor market in 2014 (since our last entry cohort is 2008), and some of the older graduates are clearing their debt.

Figure 1: Comparison of HMRC and Official Student Loan Company Records

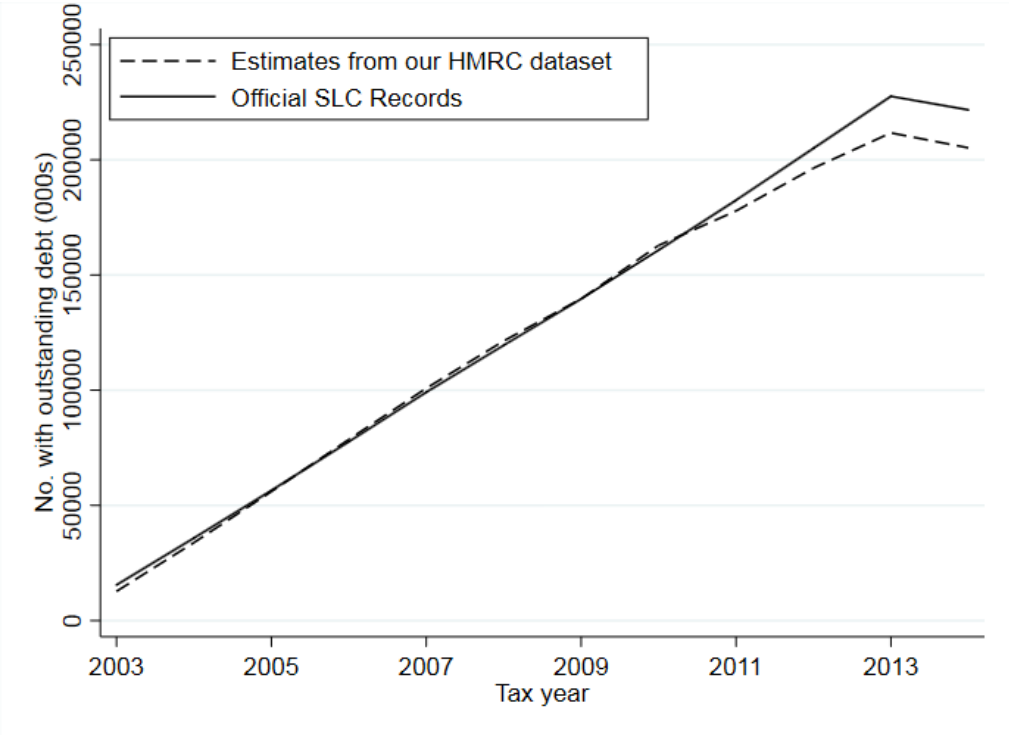

Notes: figure compares the number of borrowers in the labor market with positive debt imputed from the HMRC data with official records obtained from the SLC. The SLC numbers are divided by ten to compare directly with the HMRC numbers which are based on a 10% sample. Graduation year is approximated from the HMRC data, while the share with outstanding debt is also approximated based on voluntary repayments, income and amount borrowed.
Source: https://www.slc.co.uk/official-statistics/student-loans-debt-and-repayment/england.aspx

We observe that our estimated numbers are very similar to the SLC official records, which is a good validation of our dataset. The trends do differ slightly towards the end of the period, which is unsurprising since there are three main reasons the share with outstanding

debt is imperfectly measured. First, we do not observe the timing of earnings or voluntary repayments during the tax year, which will affect the interest applied. Second, we do not observe repayments from people overseas. Individuals who move overseas are eligible to make repayments based on a threshold that is set specifically by the SLC for each country based on the cost of living. They are required to alert HMRC when they move abroad, and based on this Britton, Shephard and Vignoles (2018) estimate that around 4% of our sample do this at some point in the first 10 years after graduation. In some cases – for example people moving to a very high paid job in Dubai – these student loan repayments from overseas could be substantial, but unfortunately we do not have any data on this. Third, we do not observe voluntary repayments after 2012. This latter factor is likely to explain the majority of the divergence in trends in 2013 and 2014. However, it is not important in terms of the overall conclusions of this paper.

We also have access to a dataset consisting of non-borrowers. This is essentially the residual dataset of anyone who is in the HMRC tax records but not in the SLC database. We observe nothing about these individuals except their age, gender and earnings. Due to the large sample size, we keep only two individuals from this non-borrower data set for every individual in the borrower data, which equates to around half of the total non-borrower sample. It includes people living and working in the UK who:

- Did not go to university;
- Went to university outside of the UK;
- Studied in the UK but did not borrow (either by choice or because they were ineligible).
- Were domiciled in Scotland, Wales or Northern Ireland, went to university in the UK and borrowed from the non-English part of the SLC;

None of these groups will be problematic for our analysis – which effectively uses this "residual" dataset as a control group – except for this final group. It is not possible to separate them out, but it is possible to quantify them. Indeed, we estimate that they account for around 10% of this residual dataset.[16]

We observe the earnings of individuals in each of the tax years between 2001/02 and 2013/14. From Table 1, this means that we observe four different earnings thresholds. We will investigate bunch at each of these thresholds following methodology described in the following section.

Table 2: Sample and student loan borrowing by cohort

Cohort	No. individuals	Average borrowing (£)
1998	14,487	7,754
1999	22,621	8,692
2000	23,506	8,626
2001	23,924	8,869
2002	23,891	9,089
2003	23,972	9,452
2004	23,577	10,240
2005	25,103	11,051
2006	25,383	14,869
2007	25,352	15,141
2008	20,847	14,085

Notes: cohort is defined based on the year the degree started. Individuals are only included in our dataset if they had entered repayment (i.e. left university) by 2011/12.

Table 2 shows the number of individuals in our sample by cohort, which is defined by the year they entered the course that we observe them on. There are roughly 25,000 people in each of the later cohorts, and slightly fewer in the earlier cohorts. In 1998, which was the first year of the income contingent loan, there are only around 15,000 individuals,

[16] This draws on author's calculations based on Table A2 from Britton, Shephard and Vignoles (2018) and https://www.gov.uk/government/collections/statistics-on-higher-education-initial-participation-rates

due to slow take-up of the loans. Individuals are only included in the dataset if they have "entered repayment" by the 2011/12, which means they need to have left university (but not actually started earning above the threshold). The sample sizes are relatively small because we have a 10% random sample of the overall cohort of borrowers. Table 2 also shows the average amount of debt individuals held upon graduation. This is initially around £8,000, increasing to around £15,000 for the 2006 cohort onwards, reflecting the increase in availability of credit the aligned with the 2006 increase in tuition fees.

Table 3: Share of borrowers with positive recorded income by cohort and tax year

	Cohort										
Year	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
2001/02	54.4	46.1	44.7	37.8	46.5	36.9	21.6	12.1	9.5	7.6	7.6
2002/03	66.4	58.5	49.6	46.6	38.9	47.3	37.5	20.2	12.3	8.9	8.7
2003/04	71.8	68.6	61.1	51.4	47.9	40.4	48.9	38.8	20.8	11.9	10.7
2004/05	60.9	59.9	57.3	51.2	43.1	40.1	33.6	40.1	30.1	16.2	11.2
2005/06	73.0	73.0	71.1	66.8	60.4	50.4	44.8	37.2	43.2	32.9	20.3
2006/07	76.4	77.1	76.1	74.8	73.2	68.4	60.6	54.7	49.8	50.0	39.0
2007/08	78.4	79.7	79.3	79.1	78.2	75.4	72.0	63.9	56.8	51.8	52.4
2008/09	87.1	88.3	88.4	89.3	89.9	89.0	88.7	85.7	78.9	77.2	75.1
2009/10	84.7	85.5	86.2	86.7	87.2	87.5	87.3	86.2	82.3	74.6	71.1
2010/11	85.1	86.1	86.6	87.5	87.8	88.3	88.8	88.1	86.8	84.5	74.5
2011/12	85.2	86.3	86.4	87.6	88.0	88.3	89.3	89.4	89.1	89.7	88.6
2012/13	84.3	85.3	85.6	86.7	87.0	87.4	88.2	88.8	88.9	89.5	89.6
2013/14	83.9	85.4	85.2	86.1	86.9	87.1	88.0	88.4	88.3	88.9	89.2

Notes: cohort is defined based on year started degree, which is typically in September of the given year. Positive recorded income is based on the tax year which runs from April 6th to April 5th.

Table 3 shows the share of borrowers that we observe with positive official earnings, by cohort and tax year. We see that for each cohort, the share of individuals with positive recorded earnings in the tax data increases with age, up to a maximum of just under 90%. Britton, Shephard, and Vignoles (2018) show that the share of borrowers in that are not in work in this dataset aligns closely with the share of graduates who are not in work in the UK

Labour Force Survey. For the younger cohorts, we see that there is a small fraction of people with positive earnings before attending university, highlighting that the sample includes some mature students. The share of individuals with positive earnings during study is also quite large at around 40% for the middle cohorts. This is likely to be primarily people who are working during study. This share is higher in the later cohorts (51.8% of the 2007 cohort in 2007/08; 75.1% of the 2008 cohort in 2008/09), and is likely driven in part by the condition that we only observe people who have entered repayment by 2011/12, although it could also reflect a rise in working during studying.

Part IV: Empirical Strategy

Our goal is to assess whether the implicit tax imposed by the income-contingent repayment distorts subsequent earnings of college graduates who receive the loans. Accomplishing this goal is complicated by the fact that the receipt of the loan directly impacts earnings. Those who attend college will have loans while those who don't attend won't have loans, and clearly the earnings of the former group will be higher.

We therefore consider two complementary empirical strategies in an effort to causally disentangle the impact of this implicit tax on earnings. The first is to exploit the kink in the repayment schedule, which is £10,000 at the start of our sample (2001/02-2004/05), transitions to £15,000 in 2005/06, before increasing to £15,795 in 2012/13 and £16,365 in 2013/14, the final year of our dataset. We follow Saez (2010) and Chetty, Friedman, and Saez (2013) in using bunching around kinks in tax schedules to identify the impact of taxes on labor supply.

Intuitively, this approach uses the deviation from a monotonic evolution of tax payers around the tax kink to identify the magnitude of response to that kink.

More specifically, (Saez 2010) shows that the elasticity of taxable income (ETI) is given by:

$$\epsilon(z) = \frac{B(\Delta z)}{k.h_0(z)\dfrac{\Delta(1-\tau)}{1-\tau_o}}$$

Here $\frac{\Delta(1-\tau)}{1-\tau_o}$ is the change in the net of tax rate, k is the kink point and $h_0(z)$ is the counterfactual density. Finally, B is the excess mass driven by the bunching and is given by:

$$B = \int_{z^*-\delta}^{z^*+\delta} h(z)dz - \int_{z^*-2\delta}^{z^*-\delta} h(z)dz - \int_{z^*+\delta}^{z^*+2\delta} h(z)dz$$

We follow Chetty, Friedman, and Saez (2013) and estimate $h_0(z)$ by fitting a flexible polynomial to the observed distribution of taxable income, excluding the area around the kink point.

This approach has the advantage that it is a fairly transparent way to assess whether tax kinks distort behavior, and it relies only on the assumption of a monotonic distribution of taxpayers around the kink. But it has the disadvantage that it ignores other reasons why individuals may bunch at that kink (especially when kinks are round numbers like £10,000 or £15,000). In the analysis we attempt to address this point through a series of placebo tests.

In particular, as noted above, we have a control sample to accompany the loan holder sample where we expect to see a labor supply response: those who did not take a loan. This group is taken from our residual tax dataset and therefore includes people who went to college but did not borrow, but the majority is made up of people who did not go to college (we are unable to distinguish between the two groups). This sample will capture any time series reasons

why there is bunching at these thresholds – but only if those reasons operate both for college-bound and non-college bound students.

Given the imperfect nature of these controls to capturing other reasons for a spike in earnings at these round numbers, we consider a complementary "difference-in-difference-in-difference" (DDD) identification strategy based on the change in the threshold from £10,000 to £15,000 in 2005/06. Before this change, individuals made loan repayments on their earnings above £10,000, leading individuals to reduce their earnings to below £10,000 and to reduce their effort above that level. But after this change, individuals earning £10,000 to £15,000 have no more loan-based incentive to lower their earnings. As a result, if there is a distortionary effect of the implicit tax on labor supply, then when the threshold is increased it should lead to more earnings from (somewhat below) £10,000 to £15,000.

Of course, a problem with this approach is that other things may have changed in 2005/06 that impact the earnings distribution of recent college graduates, such as changes to demand for skilled labor. To address this point, we use the control sample of individuals who did not receive student loans. This is not a perfect control, as it is possible that time series shocks could differ between college grads and non-grads. Moreover, this approach provides an estimate of the short run response to a change in the policy, rather than steady-state labor supply distortions estimated by kinks. But the results from this DDD analysis can provide an important confirmation (or refutation) of the kink-based approach.

Part V: Results

Graphical Evidence on Bunching

We begin with a graphical description of bunching at the implicit tax thresholds put in place by the repayment system. In Figure 2, we show the distribution of taxpayers in years when the £10,000 threshold was in place (2001/02-2004/05). Figure 3 does the same for the years when the £15,000 threshold was in place (2005/06-2011/12). In every case in our results, we refer to years by the initial year of the school year – e.g. the 2002-2003 school year is referred to as 2002.

*Figure 2: Earnings distribution of **borrowers** around the £10k threshold, 2001/02-2004/05*

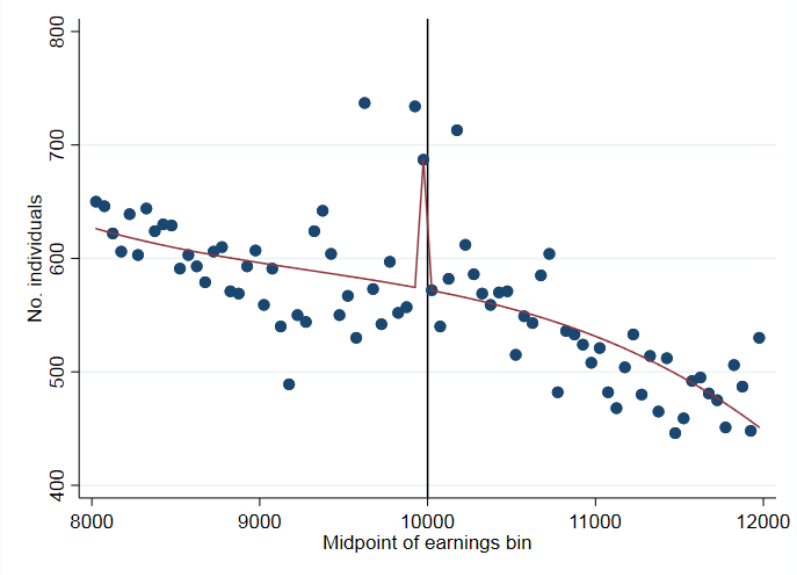

Notes: taken from linked SLC-HMRC dataset, dots indicate number of individuals within £50 earnings bins

*Figure 3: Earnings distribution of **borrowers** around the £15k threshold, 2005/06-2011/12*

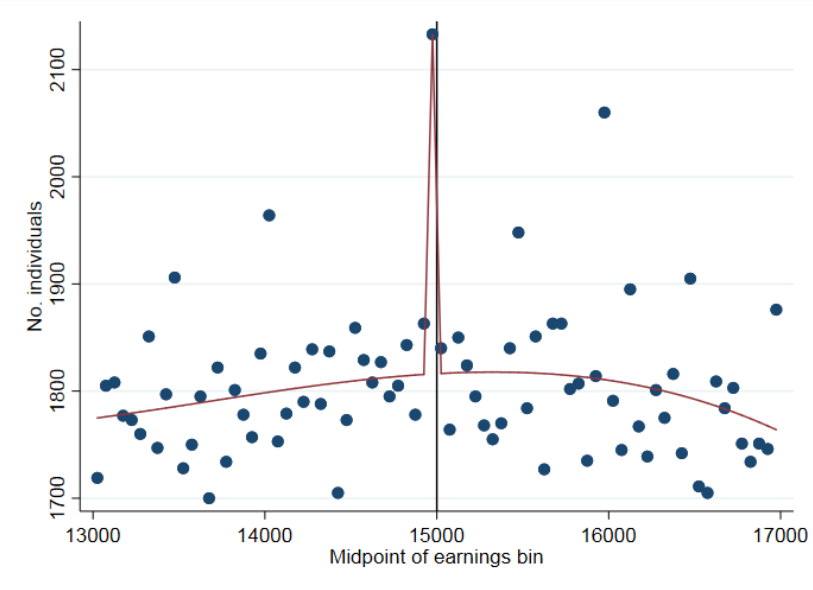

Notes: taken from SLC-HMRC dataset, dots indicate number of individuals within £50 earnings bins

These figures show some evidence of bunching, with a notable spike at the £10,000 and £15,000 notches. But this spike does not appear to be very large relative to the underlying variation in the data. These results do not immediately suggest an important effect of the thresholds on earnings.

Moreover, we see comparable spikes in the control sample of individuals without student loans, graphed in Figures 4 and 5. This sample does not have the income-contingent loan repayment incentive to bunch earnings. Yet these figures are largely indistinguishable from our main sample results, with a small spike at £10,000 and £15,000 in the relevant years.

*Figure 4: Earnings distribution of **non-borrowers** around the £10k threshold, 2001/02-2004/05*

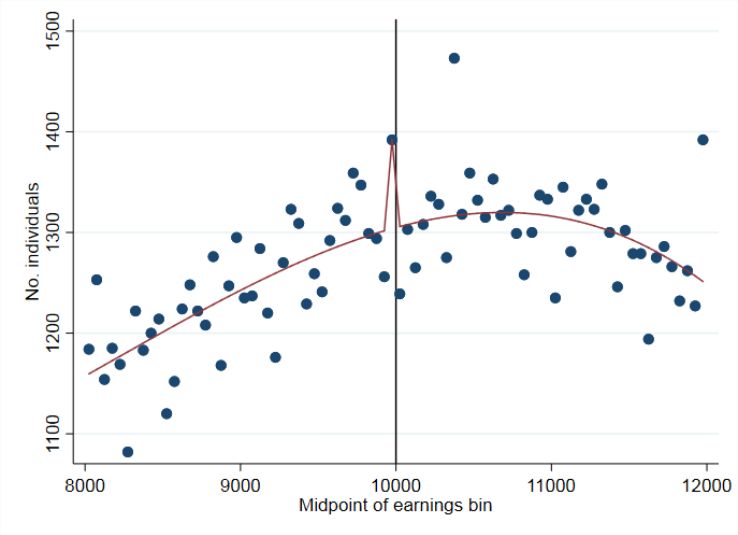

Notes: taken from "residual" HMRC dataset that consists of people who do not link to the SLC records (with the age distribution recreated). Dots indicate number of individuals within £50 earnings bins

*Figure 5: Earnings distribution of **non-borrowers** around the £15k threshold, 2005/06-2011/12*

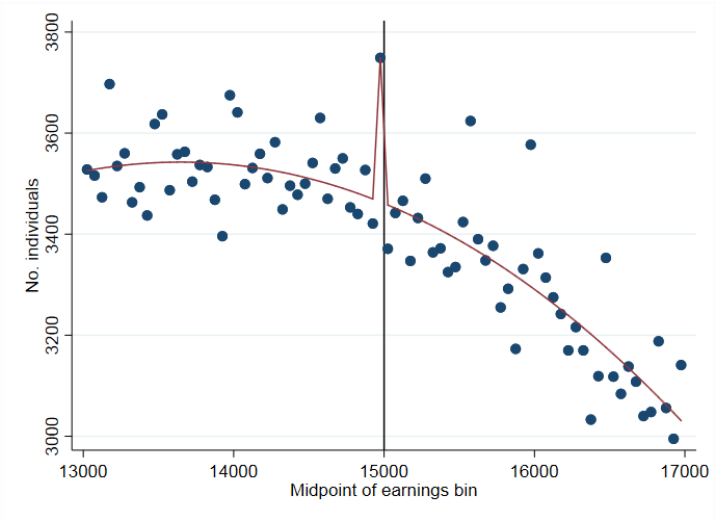

Notes: taken from "residual" HMRC dataset that consists of people who do not link to the SLC records (with the age distribution recreated). Dots indicate number of individuals within £50 earnings bins

Finally, Figures 6 and 7 show the results for the 2013 and 2014 years, when the spikes were not at round numbers. In both cases, we see no evidence of a spike – further confirming that the appearance of spikes in earlier years may be due to the round numbers used for those

thresholds. Of course, it is also possible that the weaker responses in these later years reflects delays in adjusting labor supply responses; we address this point below.

*Figure 6: Earnings distribution of **borrowers** around the £15,795 threshold in 2012/13*

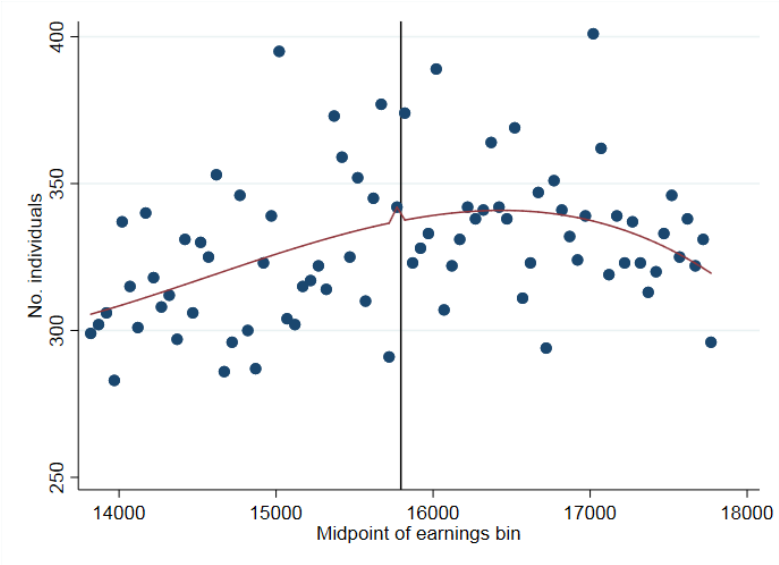

Notes: taken from SLC-HMRC dataset, dots indicate number of individuals within £50 earnings bins

*Figure 7: Earnings distribution of **borrowers** around the £16,365 threshold in 2013/14*

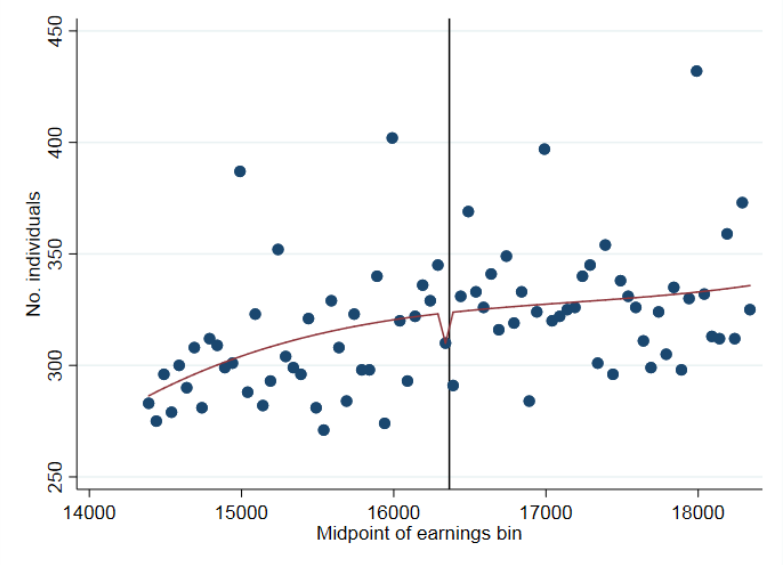

Notes: taken from SLC-HMRC dataset, dots indicate number of individuals within £50 earnings bins

Elasticities of Taxable Income

Table 4 presents our basic estimates of the elasticity of taxable income (ETI). The first column shows the results for our base sample. In the first row, we pool all years of data, testing for bunching at the relevant threshold in that year. We also show the results separately for each threshold, and show results combined for the 10k and 15k thresholds.

Table 4: ETIs calculated from the different student loan thresholds

	Borrowers	Non-borrowers	Placebo
£10k threshold	.0362***	-0.0048	0.0041
	(0.0086)	(0.0058)	(0.0044)
	57,014	*127,118*	*238,895*
£15k threshold	0.0032	0.0008	0.0028
	(0.0034)	(0.0023)	(0.0043)
	182,090	*342,052*	*97,567*
£15.8k threshold	-0.0029	0.0012	-0.0029
	(0.007)	(0.0056)	(0.0067)
	33,212	*50,066*	*33,212*
£16.4k threshold	0.0108	0.0014	.0108*
	(0.0071)	(0.0055)	(0.0065)
	31,846	*48,279*	*31,846*

Notes: Standard errors in the parentheses. *** indicates significant at the 1% level; ** 5%; * 10%.

The results across all rows in Column 1 deliver a clear message: there is little impact of the threshold on earnings. For the 15k, 15.8k and 16.4k thresholds, the estimates are small and insignificantly different to zero. For the 10k threshold years we estimate a highly significant elasticity of 0.036. However, this is very small compared to the literature: for example Adam et al. (2017) estimate elasticities of around ten times this figure for higher rate taxpayers in the UK, while the wider literature suggests central estimates of around 0.25 (see Saez, Slemrod, and Giertz, 2012, for a review).

The second column of the table shows the results for the control sample of individuals who do not have student loans. We find no meaningful effect in any year for this control sample.

Finally, in the last column, we present a placebo test for each of the individual thresholds. This is done in years where the threshold does not apply. For example, for the £15k threshold we look at bunching at the £15k threshold in the tax years 2001/02-2004/05 (when the threshold was £10k) and 2012/13-2013/14 (when the threshold was above £15k) combined. In all four of the placebo tests we get very small estimates, which are robust the specific subset of years used.

Adjustment Problems?

There are a number of possible explanations for our small results. One is that workers are slow to adjust their earnings to respond to these threshold incentives. We can test this directly by examining how the response to the new 15k threshold varied over time.

Figure 4 shows the results year by year for the 15k threshold. If the cause of our small response were adjustment problems, we would expect the spike at 15k to grow over time. But this is clearly not the case; the spike grows, then shrinks, then grows again, but in all case remains small relative to the underlying noise in the data. Our results do not appear to result from slow adjustment.

Figure 8: Earnings distribution of borrowers around the £15k threshold, 2005/06-2008/09

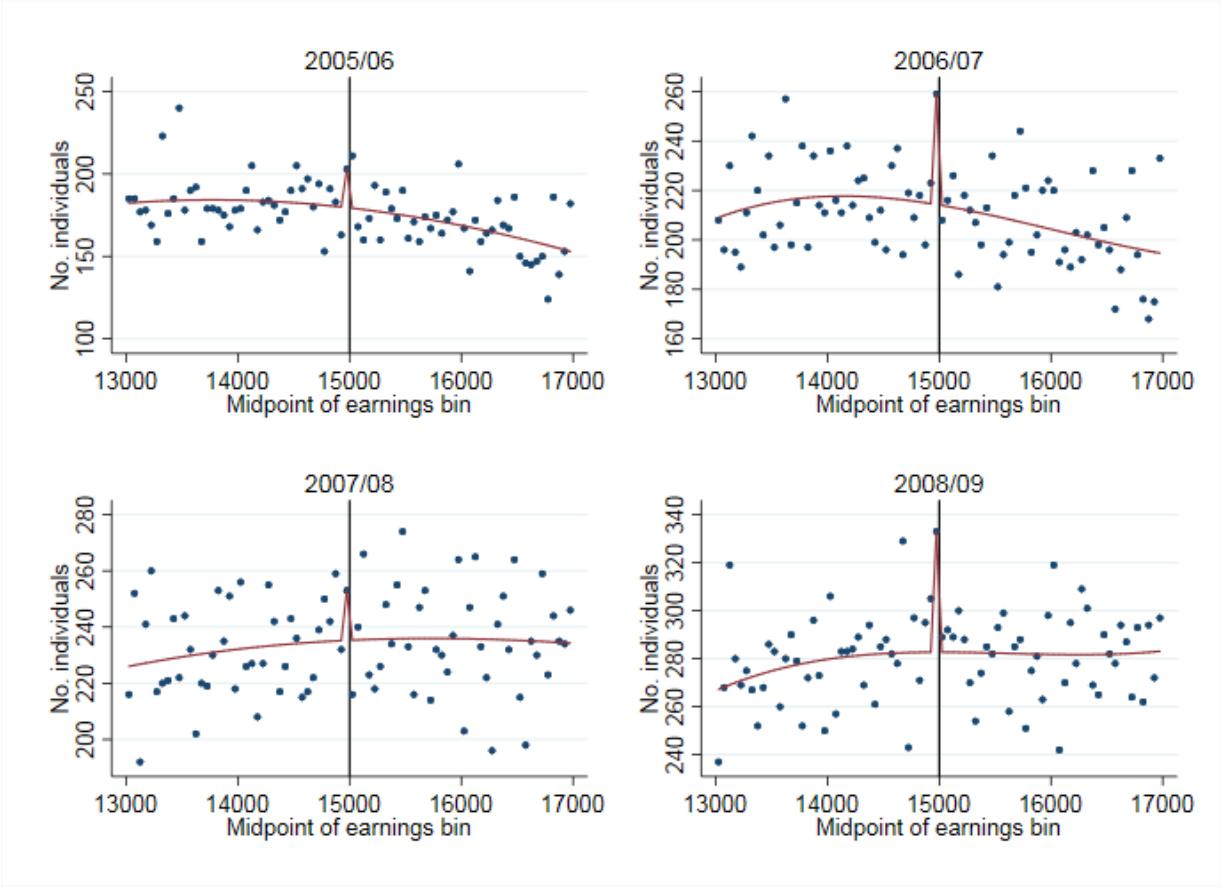

Notes: taken from SLC-HMRC dataset, dots indicate number of individuals within £50 earnings bins

Self-Employed vs. Employees

Previous work on responses to tax incentives have emphasized the role of the self-employed (for example, Adam et al. 2017), who have more scope to adjust their earnings in response to such incentives. We can explore this in our data as well by investigating bunching at the student loan thresholds amongst those who file Self Assessed (SA) tax forms. For individuals who submit SA forms, the vast majority of individuals are self-employed, for whom it is a legal requirement to submit a SA form. When looking at these individuals, we include all income from the SA form to capture other sources of income self-employed individuals may

have. This is appropriate, as student loan obligations are based on income rather than earnings (the two are equivalent for PAYE filers).

Table 5 shows results breaking down our estimates to look at self-employed individuals, distinct from employees. Given sample size restrictions, we are only able to do this for the £15k threshold due to small sample sizes. For the £15.7k and £16.3k thresholds we only have one year of data, which turns out to be insufficient. For the £10k threshold, although we have more years, there is still not enough data as there are fewer cohorts from our data in the labor market, and those in the labor market are younger, at which point the likelihood of filing an SA tax form is much lower (see, for example, Britton, Shephard and Vignoles, 2018).

The first column of Table 5 replicates the basic results from Table 4. The second column shows those same results for the PAYE (employed) sample, which excludes everybody who has a SA form, and in the final column we show results for the SA (self-employed) sample. The latter does include individuals who do not report any self-employed income but are filing a SA for other reasons (e.g. due to accruing profits from a partnership). It is not clear whether it is best to include or exclude this group, although in practice the results are unaffected by this choice.

Not surprisingly, the results for the PAYE sample mimic the overall findings. But what is striking is the lack of response among the self-employed sample – even for this group that has been found to be responsive to tax incentives, we see little impact on their earnings. The effects are larger than for PAYE, and close to significant. But even the upper bound of the confidence interval is consistent with a quite small elasticity for this population. Thus, no matter how we examine the data, we see no evidence of bunching in response to the thresholds used for student loan payments.

Table 5: Elasticity estimates by tax filing status

	All Borrowers	Employees	Self-Employed
£15k threshold	0.0032	0.0021	0.0204
	(0.0034)	(0.0034)	(0.0133)
	182,090	*170,922*	*11,168*

Notes: Standard errors in the parentheses. *** indicates significant at the 1% level; ** 5%; * 10%. Employee type is based on the type of tax form being filed (PAYE for employees; Self-Assessed for the self-employed).

Difference-in-Difference-in-Difference Strategy

One reason for the modest response estimated thus far is that individuals may be unable to exactly tie their earnings to the tax thresholds, even if they are self-employed. Therefore, we pursue a second approach, which looks at whether the change in the threshold from £10k to £15k affected income growth for people with income near to the threshold before the change.

We look at income growth from before 2005/06 to after 2005/06, the year when the threshold was first increased from £10k to £15k. We define the following three groups:

- Group 1: Pre-treatment earnings in £5k-£10k range
- Group 2: Pre-treatment earnings in £10k-£15k range
- Group 3: Pre-treatment earnings in £15k-£20k range

If the student loan program distorts income, theory suggests that we might expect the threshold increase to affect income growth differentially for the three groups: those just below the £10k threshold before the policy change (Group 1) experience no income effect as a result of the policy, but do experience a substitution effect, suggesting they might increase their income by more as a result of the reform, as less of their income growth is spent on loan repayments.

On the other hand, those just above the new £15k threshold (Group 3) experience no substitution effect, as their marginal income is taxed at a rate of 9% as before, but they do experience an income effect, gaining 9% of £5k (£450). This pure income effect might induce those in the £15k-£20k range to increase their earnings by less as a result of the reform. For those in the £10k-£15k range (Group 2), the effect is ambiguous as they experience a mix of income and substitution effects.

We investigate these theoretical predictions using a triple differences (DDD) framework. For our "treated" group, we use the earnings growth of borrowers from before the threshold change to after it (i.e. between 2004/05 and 2005/06). For our control group, we use earnings between 2003/04 and 2004/05, which should be unaffected by the reform. For this we use a different (older) cohort so we are observing the treatment and control groups at the same point in their careers. For example, if our treated group is the 1999 cohort, we would look at their earnings growth between 2004/05 and 2005/06, which correspond (in the majority of cases) to three and four years after graduation. Our control group in this case would be the 1998 cohort, looking at their earnings growth between 2003/04 and 2004/05, also three and four years after graduation.

However, this does not account for wider macroeconomic changes that might affect earnings growth at the same point in the career of two different cohorts differently. To deal with this, we use our non-graduate sample as an additional control group. Specifically, we subtract the difference in earnings growth for the 1999 and 1998 cohorts of non-borrowers (where "cohort" imputed based on age). Since neither of these cohorts is affected by the reform, removing this difference should net out the wider macro shocks.

We also follow a matching technique suggested in Abadie, Diamond, and Hainmueller (2010) and implemented more recently in Gruber, Jensen, and Kleven (2017). We take everybody in our treated sample (in this case, the 1999 cohort) and match them to their "nearest neighbor" in the control sample.[17] Specifically, for each individual, i, we select individual j to satisfy:

$$\min_j \{\sum_{t=1}^{3} |x_{it} - x_{jt}|\}$$

We put equal weight on each of the pre-reform earnings. We do this with replacement, although the findings are robust to this decision. We use this "synthetic" group as our control for both the borrowers and non-borrower samples in the DDD regressions.

The graphical representation of our DDD approach using the synthetic control group is given in Figure 9 for Groups 1, 2 and 3 as defined above. In each case, the solid colored line is average earnings for the treated borrower sample, while the colored dashed line is average earnings for the untreated borrower sample. The gray lines are the equivalent cohorts to the treated and untreated cohorts in the non-borrower sample. We see that in both the borrower and the non-borrower samples, the pre-treatment trends are almost perfectly on top of each other, suggesting our matching approach works very well. A simple version of the DDD estimation would take the difference between the colored solid and dashed curved and subtract the difference between the gray solid and dashed curves (for group 1, this would result

[17] We find in practice that the pre-treatment trends for the treatment and control groups are not perfectly parallel (see Figure A1 in the Appendix).

in a very small positive effect of the reform, with the non-borrower group making very little difference).

However, Table 6 shows the results from the DDD within the proper framework, baed on the following regression model:

$$y_{it} = \beta_0 + \beta_1 t + \beta_2(t * B_i) + \beta_3 T_i + \beta_4(T_i * B_i) + \beta_5 B_i + \beta_6 Post + \beta_7(Post * B_i) + \beta_8(Post * T_i) + \beta_9(Post * T_i * B_i) + \epsilon_{it} \quad (1)$$

Here y_{it} is income for individual i at time t, t is a linear time trend, T is a dummy indicating that the individual is in the treated cohort, B is a dummy for being in the borrower sample and Post is a post reform dummy (in Figure 9, this is in the fourth year). So in this framework we control for differences across treatment vs. control cohorts (β_3), borrowers vs. non-borrower (β_5), post vs. pre-reform (β_6).

Figure 9: Income growth around the threshold increase for treated and untreated groups, for borrowers and non-borrowers

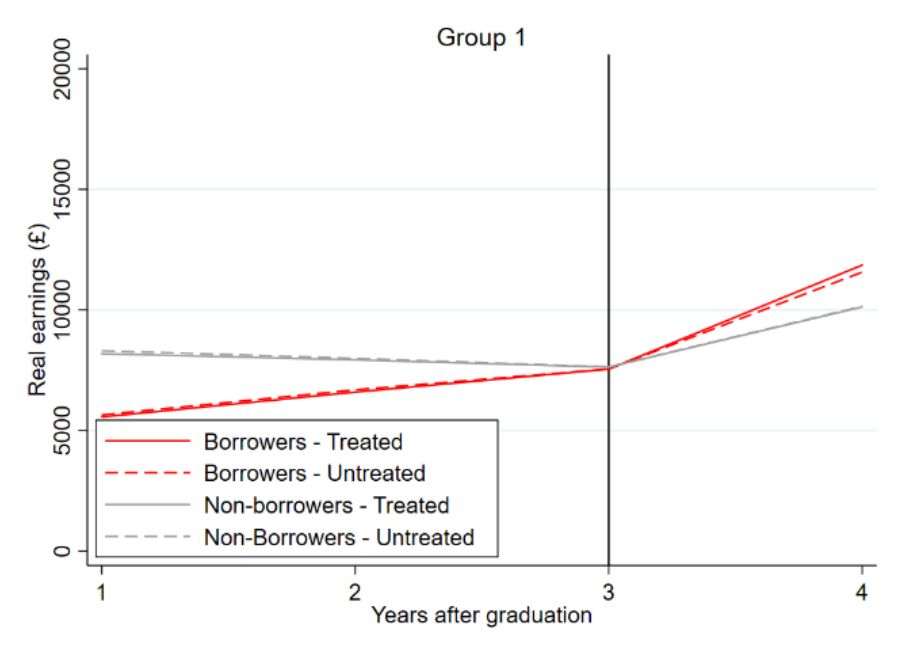

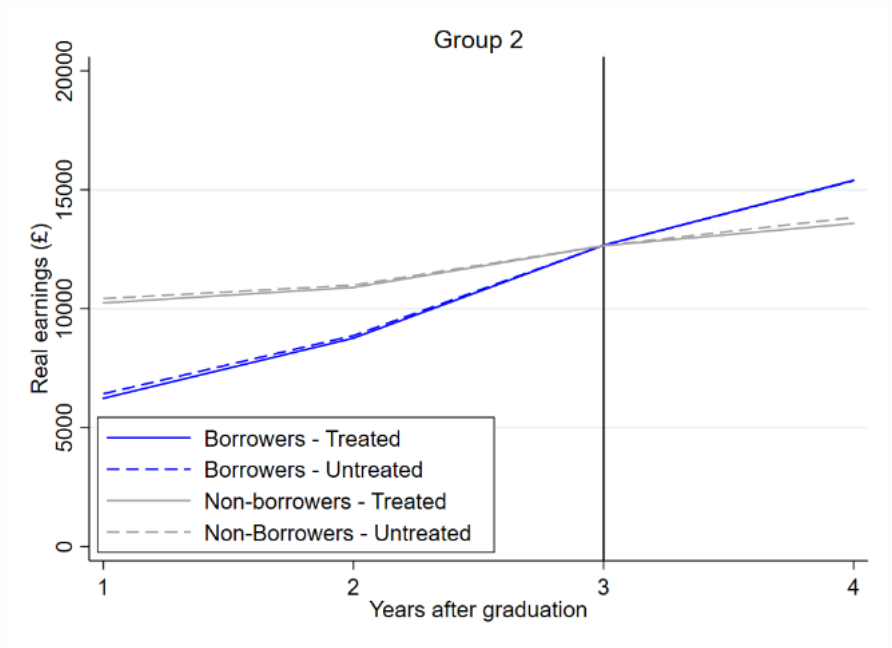

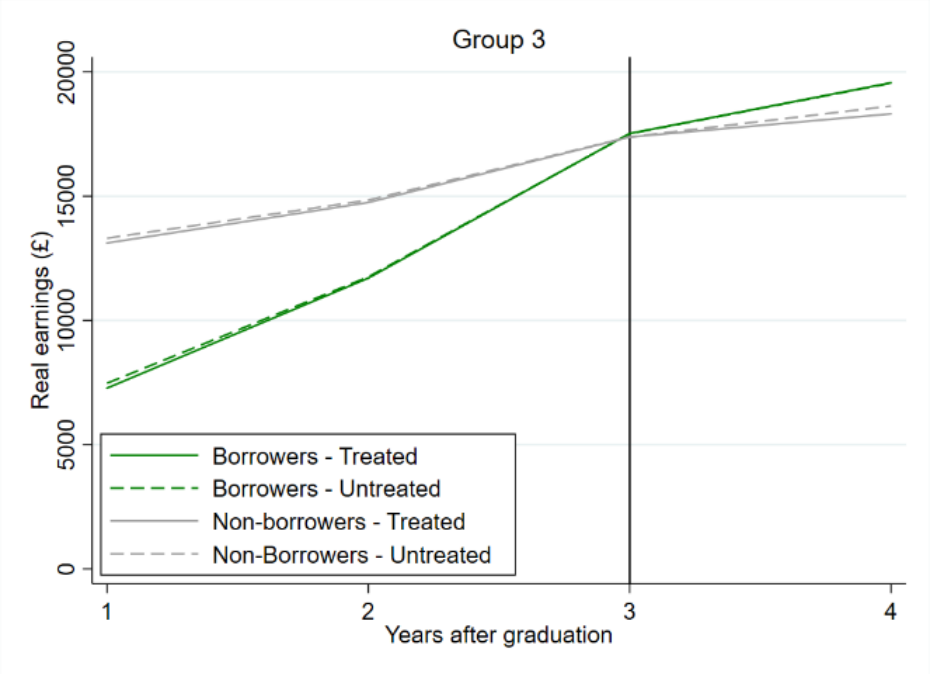

Notes: Figure shows average earnings before and after the increase in the repayment threshold from £10k to £15k in 2005/06. The treatment group is the 1999 cohort, for whom t=4 is 2005/06, the first post-treatment year. The untreated group is the synthetic 1998 cohort, for whom t=4 is 2004/05, when the threshold is still £10k. The "treated" and "non-treated" non-borrowers are not affected by the reform, but the cohorts are equivalent to the borrower cohorts. Group 1 is those with earnings in the £5k-£10k range in the pre-reform years, Group 2 is £10-£15k and Group 3 is £15k-£20k. Figures condition on having non-zero earnings.

We also control for "second-level" differences between borrowers in treatment cohorts vs. borrowers in non-treatment cohorts (β_4), borrowers post-reform vs. pre-reform, (β_7), and treatment cohorts post-reform vs. pre-reform (β_8). The coefficient of interest therefore is (β_9) which is the estimated effect of the reform on the earnings growth of the treatment group, relative to the control group, among borrowers, relative to non-borrowers, from before to after the reform.

Table 6: DDD regressions investigating the impact of the £10-£15k threshold increase on earnings growth

	[1]	[2]	[3]
Group 1 (5-10k)	251.4	-226.5	140.4
	(263.7)	(235.1)	(222.9)
	25,713	*25,625*	*21,706*
Group 2 (10-15k)	257.6	169.6	222.4
	(203.2)	(178.6)	(206.5)
	43,547	*37,826*	*22,205*
Group 3 (15-20k)	309.2	12.2	-20.74
	(194.4)	(220.6)	(314.6)
	44,711	*28,862*	*13,700*
Treated cohort	1999	2000	2001
Untreated cohort	1998	1999	2000
Pre-reform years	3	2	1

Notes: Group 1, 2 and 3 are defined based on pre-reform income. All estimates here are from separate regressions, with standard errors given in the parentheses. * = significant at the 10% level, ** = 5%, ***=1%. DDD indicates a triple-diff regression.

Table 6 shows our estimated effect of the reform on earnings growth for each of the three groups. Within each column the estimates are from separate regressions, with the corresponding sample sizes given below. Column 1 shows the estimated effects from the DDD model based on the data plotted in Figure 9, where the treatment group is the 1999 cohort and

the control group is the synthetic control group constructed from the 1998 cohort. For this regression, there are three pre-reform years. The estimates show insignificant effects for all three groups.

Column 2 and 3 of Table 6 then use different samples to try to answer the same question. Column 2 looks at growth of the 2000 cohort between two and three years after graduation, using the 1999 cohort as a control, while column 3 uses the 2001 cohort between one and two years after graduation with the 2000 cohort as a control. Consequently the number of pre-reform years reduces in each case. In all cases we observe insignificant results - in some cases with the opposite sign - suggesting that earnings growth was not affected by the reform to the repayment threshold.

This finding is consistent with the results from our bunching estimation, which suggests that the ETI of graduates is effectively zero.

Part VI: Conclusions

A college degree is increasing a pre-requisite for success in the modern knowledge-based economy. Indeed, having a more highly educated population in a location increases the returns even for those without a college education (Moretti 2004). At the same time, inflation in college tuitions has exceeded general price inflation in many nations. In the U.S., for example, the overall CPI since 1983 has risen by less than 250% - while average tuition and fees have risen by 800%. The result is a large and growing burden of student debt that has been deemed unaffordable for many. A direct approach to address affordability is to more tightly link college debt repayments to income, as is true for a minority of U.S. loan recipients – but almost all college attendees in the U.K.

One concern with such an approach is that income-related debt repayments put in place an implicit tax – potentially in a range of incomes where the phase out of other social insurance programs already. In that case, there may be a distortion to labor supply decisions that may offset the welfare gains from insuring college graduates against income risk.

We assess whether such labor supply distortions are quantitatively important by studying the income contingent loan repayment program for college graduates in the U.K. Thresholds for repayment at £10k, later £15k, suggest that if this program is distorting earnings, there should be bunching at those thresholds. In fact, the bunching that we see at these thresholds is very modest – and is present in a control sample that should not respond to these incentives. To further confirm our results, we turn to an alternative strategy, examining whether the density of the income distribution in the range between 10k and 15k increases when the threshold is moved upwards – and we find it is not. We therefore conclude that there are not quantitatively important labor supply distortions from this policy.

The income contingent loan program effectively converts college loans into a social insurance policy. As with any such policy, there may be other margins of distortion. In particular, such a policy may induce excess college attendance among those who know that such education will have little return – but who will have low earnings later on and so therefore won't have to pay back. At the same time, there is reason to believe that individuals are sub-optimally deciding to pass on college (Oreopoulos 2009).

Another important question is how this financing mechanism impacts behavior once at college. Recent studies show that loan financing, relative to grants, can lead students to change

to more renumerative majors.[18] Would this impact be lessened if the loan repayment were tied to income? Future work could usefully explore these questions

[18](Field 2009; Rothstein and Rouse 2011).

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association* 105 (490): 493–505.

Adam, Stuart, James Browne, David Phillips, and Barra Roantree. 2017. "Frictions and Taxpayer Responses: Evidence from Bunching at Personal Tax Thresholds." *Institute for Fiscal Studies Working Paper W* 17.

Brewer, Mike, Emmanuel Saez, and Andrew Shephard. 2010. "Means-Testing and Tax Rates on Earnings." *Dimensions of Tax Design: The Mirrlees Review* 1.

Britton, J., L. Dearden, N. Shephard, and A. Vignoles. 2018. "Is Improving Access to University Enough? Socio-Economic Gaps in the Earnings of English Graduates." *Oxford Bulletin of Economics and Statistics*, no. Forthcoming.

Britton, J., Shephard, N. and Vignoles, A. 2018. "A Comparison of Sample Survey Measures of Earnings of English Graduates with Administrative Data." *Journal of the Royal Statistical Society Series A (with Discussion)* Forthcoming.

Chapman, Bruce, and Andrew Leigh. 2009. "Do Very High Tax Rates Induce Bunching? Implications for the Design of Income Contingent Loan Schemes*." *Economic Record* 85 (270): 276–289. https://doi.org/10.1111/j.1475-4932.2009.00554.x.

Chetty, Raj. 2012. "Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." *Econometrica* 80 (3): 969–1018. https://doi.org/10.3982/ECTA9043.

Chetty, Raj, John N Friedman, and Emmanuel Saez. 2013. "Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings." *American Economic Review* 103 (7): 2683–2721. https://doi.org/10.1257/aer.103.7.2683.

Eissa, Nada, and Jeffrey B. Liebman. 1996. "Labor Supply Response to the Earned Income Tax Credit." *The Quarterly Journal of Economics* 111 (2): 605–637.

Field, Erica. 2009. "Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment at NYU Law School." *American Economic Journal: Applied Economics* 1 (1): 1–21.

Friedman, Milton, and Simon Kuznets. 1945. *Income from Independent Professional*. New York: National Bureau of Economic Research.

Gruber, Jon, and Emmanuel Saez. 2002. "The Elasticity of Taxable Income: Evidence and Implications." *Journal of Public Economics* 84 (1): 1–32.

Gruber, Jonathan, Amalie Jensen, and Henrik Kleven. 2017. "Do People Respond to the Mortage Interest Deduction? Quasi-Experimental Evidence from Denmark." National Bureau of Economic Research.

Hout, Michael. 2012. "Social and Economic Returns to College Education in the United States." *Annual Review of Sociology* 38: 379–400.

Kleven, Henrik J., and Mazhar Waseem. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan." *The Quarterly Journal of Economics* 128 (2): 669–723.

Moretti, Enrico. 2004. "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions." *American Economic Review* 94 (3): 656–690.

Moriya, Asako S., Thomas M. Selden, and Kosali I. Simon. 2016. "Little Change Seen in Part-Time Employment as a Result of the Affordable Care Act." *Health Affairs* 35 (1): 119–123.

Mueller, Holger M., and Constantine Yannelis. 2018. "The Rise in Student Loan Defaults in the Great Recession." *Journal of Financial Economics*.

Mulligan, Casey B. 2015. "The New Full-Time Employment Taxes." *Tax Policy and Economy* 29 (1): 89–132.

Oreopoulos, Philip. 2009. “Would More Compulsory Schooling Help Disadvantaged Youth?” *The Problems of Disadvantaged Youth: An Economic Perspective* 85.

Oreopoulos, Philip, and Uros Petronijevic. 2013. “Making College Worth It: A Review of Research on the Returns to Higher Education.” National Bureau of Economic Research.

Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva. 2014. “Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities.” *American Economic Journal: Economic Policy* 6 (1): 230–271.

Rothstein, Jesse, and Cecilia Elena Rouse. 2011. “Constrained after College: Student Loans and Early-Career Occupational Choices.” *Journal of Public Economics* 95 (1–2): 149–163.

Saez, Emmanuel. 2010. “Do Taxpayers Bunch at Kink Points?” *American Economic Journal: Economic Policy* Vol 2 (3): 180–212. https://doi.org/10.1257/pol.2.3.180.

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz. 2012. “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review.” *Journal of Economic Literature* 50 (1): 3–50.

Shireman, Robert. 2017. “Learn Now, Pay Later: A History of Income-Contingent Student Loans in the United States.” *The ANNALS of the American Academy of Political and Social Science* 671 (1): 184–201.

Sieg, Holger, and Yu Wang. 2018. “The Impact of Student Debt on Education, Career, and Marriage Choices of Female Lawyers.” *European Economic Review* 109: 124–147.