

# Grades across universities over time

GERAINT JOHNES<sup>i</sup>  
Lancaster University

KWOK TONG SOO<sup>ii</sup>  
Lancaster University

January 2014

## Abstract

This paper examines the determinants of degree outcomes in a sample of UK universities from 2005 to 2012. We use stochastic frontier methods to account for differences in efficiency across universities and over time. The quality of the student intake, the university's research performance, and student satisfaction are the main determinants of degree outcomes. There is no evidence of grade inflation. Decomposing the determinants of degree outcomes, we find that good entry grades yield a higher return in traditional universities than in new universities. Although high quality universities appear to be more lenient, we find little evidence that universities of different quality differ in their propensity to inflate grades over time.

JEL Classification: I21, I23.

Keywords: Degree outcomes; UK universities; stochastic frontier analysis; Blinder-Oaxaca decomposition.

---

<sup>i</sup> Department of Economics, Lancaster University Management School, Lancaster LA1 4YX, United Kingdom. Tel: +44(0)1524 594215. Email: [g.johnes@lancaster.ac.uk](mailto:g.johnes@lancaster.ac.uk)

<sup>ii</sup> Department of Economics, Lancaster University Management School, Lancaster LA1 4YX, United Kingdom. Tel: +44(0)1524 594418. Email: [k.soo@lancaster.ac.uk](mailto:k.soo@lancaster.ac.uk)

## INTRODUCTION

*“Everyone at this graduation is first in his class. What a grand and glorious accomplishment! The diplomas are all here and since every one of you is first in his class, you can simply come up and get them in any old order you please”.*

Lao She, *Cat Country*.

In any sphere of activity where standards are set, assessors face the challenge of ensuring the constancy of the level of performance that is required to meet the given standard. In certain contexts this is a straightforward task, but in others, where objective measurements are difficult or when the context within which the evaluation is taking place is changing, assessment to a consistent standard is more difficult. In such circumstances, incentives may be present that serve to discourage assessors from being consistent, and the outcome will result from a tension between these incentives and forces working in the opposite direction; this is notably the case where success in meeting the standard might be used as a performance indicator to which reward is attached. In the field of monetary policy, for example, such conditions might lead to price inflation. In the field of education, which is the focus of the present paper, they might lead to grade inflation.

Grade inflation has been a perennial source of concern for policy makers. In the United Kingdom, the Higher Education Academy has recently launched its ‘Transforming Assessment in Higher Education’ pilot scheme (<http://bit.ly/X06tLP>) amidst substantial press coverage.<sup>1</sup> This project aims to refine the process of student assessment, thereby compensating for the fact, noted in several studies, that the ability of the current system of degree classification to distinguish between students’ performance has declined (Universities UK, 2007, 2012; Higher Education Academy, 2012). But equally in other countries, the tendency for grade inflation to occur has been noted. In the US, Stuart Rojstaczer has written about this tendency in the Washington Post (28 January 2003), and has subsequently recorded the phenomenon on his website [www.gradeinflation.com](http://www.gradeinflation.com). He finds that, since the 1980s, average GPA scores across US colleges has risen from below 2.9 to 3.1. In related work, Rojstaczer and Healy (2010) detail how this increase in average scores varies across types of institution.

If the increase in grades serves to compress more and more students into a smaller number of grade classifications, then clearly the process of grade inflation entails a loss of information to users – which may include graduate schools and employers. This, in turn, can result in a deterioration in the quality of employer-employee matches, thereby harming economic efficiency. While grade inflation concerns a nominal variable, it can therefore have real costs.<sup>2</sup>

Yet it is far from clear that the increases in assessment scores that are observed in practice are indeed the result of grade inflation. A number of other factors may be at work, and it is desirable that we should, as far as possible, seek to separate these out. For example, technological advances and improvements in quality assurance procedures may have led to pedagogical improvements in universities that enable them to add more value to their students than was previously the case. Likewise, similar improvements in primary and secondary education may mean that universities can build on students that enter higher education with stronger foundations of knowledge and understanding. Moreover, changes in the efficiency with which students sort themselves across universities may further contribute to rising performance over time. Our goal in this paper is to shed as much light as possible on the factors that underpin the apparent improvements in students' performance over time, and hence to establish the extent to which this is attributable to grade inflation.

The remainder of the paper is structured as follows. In the next section, we present a brief review of the relevant literature. This is followed by a consideration of our methodological approach and the data sources used in the empirical analysis. We then report the results of this analysis, and conclude with a restatement of our findings and some suggestions for future research.

## **I. LITERATURE REVIEW**

Marks (2002) argues that norms are public goods. These norms include those that attach to the evaluation of work that students submit for assessment – they apply across whole cohorts of students and are thus non-rival and non-excludable. In common with other public goods, there are incentives for consumers to free ride. In the context of norms, people are generally

happy for others to subscribe to standards, but if adherence is costly they are happy to seek ways of circumventing them themselves. Higher education institutions aspiring to be classed among the elite may wish to set high standards in order to enhance their prestige, but at the same time yield to the temptation to manipulate metrics by which they are judged.

McKenzie and Tullock (1981) provide an interesting rationale for this in their analysis of how tuition fees in the US rose during the rapid post-Sputnik expansion of higher education. Institutions that subsequently found difficulty in recruiting sufficient students were unable to reduce tuition fees since their costs had risen in line with their revenues; they had therefore to find a different way of pricing themselves into the market, and increasingly lenient grading practices may have been used to achieve this. Freeman (1999) has tested this idea by examining the relationship between grading standards and labour market prospects in different disciplines. His findings support the hypothesis that grading is more lenient in disciplines where the price of higher education needs to be relatively low in order to compensate for a weak labour market for graduates. Other studies that report that average grades have been rising at different rates in different subject areas include Sabot and Wakeman-Linn (1991), Anglin and Meng (2000), and Jewell et al. (2013). The first of these is particularly interesting since it provides evidence for a ‘demand curve’ where the strictness of grading practice is regarded as a price.

The tension between the public good nature of norms and the incentive to cheat on those norms has led several authors to view grade inflation in a game theoretic context. Such studies include those by Correa (2001) and Johnes (2004). The latter considers a simple model in which there are two students, each of whose utility is determined by their rank order in grading. So the  $i$ th student enjoys utility of  $U_i = x_i - x_j, j \neq i$ , where  $x_i$  denotes the  $i$ th student’s grade. Suppose that  $x_i$  depends on some absolute measure of the value of  $i$ ’s work,  $x_{0i}$ , but that it is also influenced by pressure that the students bring to bear on the assessor, so that  $x_i = x_{0i} + \delta_i + \sigma\delta_j, j \neq i$ , where  $0 < \sigma < 1$ . Here the  $\delta$  terms are binary variables reflecting student-specific pressure, and students each choose the value of this variable to maximise their utility. It is readily seen that the Nash equilibrium is one in which each student sets this equal to one, thus creating pressure for grade inflation.

Chan et al. (2007) have developed the idea of grade inflation as a game further, by considering the schools to be the players. Schools signal the quality of their graduates to employers using the grades that they award. This signal is made up in part of a component that reflects students' abilities but in part also of a component that is due to schools wanting to help their own students in the face of competition from other schools. Employers know this, but they do not know how much weight is assigned to each component. It is therefore in schools' interest to inflate grades.

In the same spirit, work by Achen and Thurnau (2009) investigates pressures that university departments face vis-à-vis each other. Where modules form part of the compulsory core of a degree programme, they expect to find relatively low grades; where, on the other hand, the modules are optional, and thus need to compete with other modules for student registrations, they expect grading to be more lenient. Where assessment relies on subjective judgements, grading is expected to be more lenient, because assessors seek thereby to avoid the costs imposed by students who try to query low grades. Finally, departments that are struggling to recruit students are likely to offer more lenient grading. Some support for each of these hypotheses is found using data from the University of Michigan.

The last of Achen and Thurnau's hypotheses might, by extension, be viewed as implying that more prestigious universities are, given their strong position in the market, less likely than others to succumb to the temptations of grade inflation. Yet Popov and Bernhardt (2013) argue otherwise. They develop a model in which good universities respond to a rising demand for highly skilled workers by relaxing the conditions under which they award top grades; they do this because they are seeking to optimise employment outcomes for their graduates, and they can exploit the fact that employers are not able to distinguish between relatively strong and weak students within the top grade. In support of their theory, Popov and Bernhardt provide evidence on grade inflation over a forty year period in a range of US institutions.

Several developments in the technology of educational production may have served to improve performance. For example, Bracey (1994) suggests that the introduction of word-processing software allowed students to improve the presentation of their work. Innovations of this kind are sporadic and one-shot, however, and do not do much to account for a consistent trend over a prolonged period. Nevertheless, technology more generally might

evolve gradually in such a way that it contributes to increasing achievement over time, and this raises the question of how best to capture such effects in empirical analyses.

Efforts to evaluate the extent of grade inflation in the US include contributions by Kolevzon (1981) and Compton and Metheny (2000). Both claim to find evidence that grade inflation is indeed present. In the UK there has been some evidence of grade inflation at secondary school level. Exercises that have presented cohorts of students, years apart, with identical problems suggest that while some of the apparent gain in performance is genuine, some has also been illusory (Tymms and Fitz-Gibbon, 2001).

Evidence of grade inflation at the level of higher education has been harder to come by, however. Johnes and McNabb (2002) and Johnes (2004) adopt an educational production function approach, using stochastic frontier analysis to model a measure of the university average of students' grades, using a panel of universities as the data set. The explanatory variables include average student scholastic achievement on entry to university and a variety of other indicators of student demographics. Changes in these variables are likely to account in part for changes, over time, in university average grades. Likewise, changes in the efficiency with which universities convert inputs to outputs, measured by the non-normal component of the residual in the stochastic frontier estimator, likely accounts for some of the change in grades over time. The model specification also includes a full set of time dummies, and the authors argue that changes in the coefficients on these dummies reflect grade inflation – since they cannot be due to changes in input quality or in efficiency. The results suggest that there was some grade inflation over a short period during the mid-1980s, but not in the 1990s.

The empirical work reported in the remainder of this paper updates and builds on the Johnes and McNabb approach. In particular, it allows us to investigate whether there has been any further grade inflation in the UK higher education system since the early part of the last decade. In addition, we are able to perform a comparison between the conflicting findings of Achen and Thurnau (2009) and Popov and Bernhardt (2013) on the relationship between university quality and grade inflation.

## II. DATA AND METHODS

The data used in the empirical analysis that follows come from the Sunday Times University Guide, published annually in September each year. For each higher education institution in the UK they provide information on a variety of variables. These variables include the student-staff ratio, library expenditure, expenditure on other facilities, research performance (as measured by the Research Assessment Exercise), and student satisfaction (measured by the proportion of students expressing satisfaction on the overall measure in the National Student Survey). The data set also contains information about the average performance of students in each institution both on entry (measured by the mean tariff score achieved in the upper secondary school exit qualifications, typically A levels) and on exit (measured by the proportion of students graduating with first or upper second class honours degrees; henceforth, good degrees). All monetary values have been converted into real terms using the consumer price index with base year in 2005. Descriptive statistics for these variables are reported in Table 1. Figure 1 shows a clear rising trend in the percentage of good degrees in our sample at all points of the distribution (with the possible exception of the upper tail).

[Table 1 here]

[Figure 1 here]

We shall be interested in establishing the relationship between university output – as measured by the proportion of students graduating with good degrees – and input. In order to ensure that the appropriate input variables are used, we study output in a given year as a function of, *inter alia*, the quality of students recruited three years earlier, reflecting the duration of a typical bachelors programme in the UK.<sup>3</sup> With the National Student Survey starting in 2005, the data available allow us to analyse degree results from 2004-05 through 2011-12.

We follow the lead of Johnes and McNabb (2002) and Johnes (2004) in modelling the production of good degrees in a stochastic frontier framework. This allows us to evaluate the effects of (i) changing inputs, including changes in the quality of the student intake over time, (ii) changing efficiency in the conversion of inputs to outputs, reflecting genuine changes in productivity in the higher education sector, and (iii) through the use of time dummies, grade

inflation. By using time dummies to capture grade inflation, we are defining grade inflation to be any change in grades that is not explained by changes in the inputs or efficiency<sup>4</sup>. The inputs used in the production function are those discussed above: entry grades, student-staff ratio, library expenditure, expenditure on other facilities, research performance, and student satisfaction. We believe these are important inputs in the production of university degrees; Smith and Naylor (2001) and Wossmann (2003) have used similar types of variables in analysing the determinants of undergraduate performance in the UK (Smith and Naylor) and international student performance in science and mathematics (Wossmann).

Stochastic frontier models were introduced into the literature by Aigner et al. (1977), and have been widely used in the economics of education as a means of evaluating cost functions (see, for example, Izadi et al., 2002; Johnes and Johnes, 2009). The stochastic frontier approach is based on the idea that no university can outperform the production frontier; as a result, barring noise, all observations lie within the frontier. To implement this, the error term in the regression model is divided into two components: the classical error term, and a one-sided term which represents inefficiency. Therefore, the stochastic frontier approach is different to conventional regression methods, in which the best-fit line passes through the centre of the data, and does not allow universities to differ in terms of their efficiency<sup>5</sup>. With the benefit of a panel data set, a variety of models become available for analysis (Greene, 2005). The results reported below are obtained using a true random effects specification in which

$$(1) \quad y_{it} = \alpha_i + \beta' \mathbf{X}_{it} + \gamma_t + v_{it} + u_{it}$$

$$(2) \quad v_{it} \sim N[0, \sigma_v^2] \text{ and } u_{it} = |U_{it}| \text{ where } U_{it} \sim N[0, \sigma_u^2].$$

The dependent variable,  $y_{it}$ , is the logged proportion of degrees awarded at first or upper second class honours level, and  $\mathbf{X}_{it}$  denotes the vector of logged explanatory variables; in each case the  $i$  and  $t$  subscripts represent university and year. The  $\gamma_t$  are time dummies; an increase in the value of coefficients of these dummies as  $t$  rises would reflect a measure of grade inflation. This specification of the model allows efficiency to be measured by the  $u_{it}$  term, but allows institution-specific variation that is due to unobserved heterogeneity other than efficiency differentials to be captured by the  $\alpha_i$ . This represents an advance over the



previous literature, which is unable to separate efficiency differences from other forms of heterogeneity across institutions (Greene, 2005).

The  $\alpha_i$  may be estimated with either fixed effects or random effects used to capture variation across universities. Although the true fixed effects model is desirable in that it allows for university-specific effects to differ across universities, we have made use of the true random effects model because the true fixed effects model suffers from an identification problem. That is, the signal-to-noise ratio ( $\lambda = \sigma_u/\sigma_v$ ) in the true fixed effects model is very large, so that the log-likelihood becomes quite flat, leading to numerical maximisation problems (see Belotti et al 2013). As a result, the parameter estimates from the true fixed effects model are not very stable. In addition, Greene (2005) shows that the predicted efficiency scores from the true fixed effects and true random effects models are very similar to each other.

The half-normal specification of the efficiency term represents a choice on the part of the analyst;  $u_{it}$  must be non-normal in order to distinguish it from  $v_{it}$ , but a range of possibilities exists, including gamma, exponential and truncated normal. In our experience the choice makes little difference to results, and since half-normal is fairly standard in the literature we choose this option. It is worth noting that the assumption that efficiency follows a distribution that is non-normal is just that – an assumption. In the present analysis, we would expect this assumption to be satisfied as a consequence of the fierce competition that exists between institutions of higher education in the UK; this competition should lead most institutions to be relatively efficient in converting their inputs into good degree results, with a small number of less efficient outliers.

It is possible to restrict the one-sided error term to be constant across all time periods for each institution such that  $u_{it} = u_i$ . We choose not to impose this restriction, and so we allow flexibility in the model to accommodate changes in efficiency over time for each institution. The  $u_{it}$  are recovered using the method of Jondrow et al. (1982) and, given the log specification of the model, may be expressed as efficiency scores using the transformation  $\exp(-u_{it})$ .

Once the coefficients of the preferred equation are estimated, we conduct a three way Blinder-Oaxaca (1973) decomposition of the gap between groups in the proportion of good

degrees awarded *at the frontier*,  $y_1 - y_2 - (u_1 - u_2)$ , where the subscripts are group indicators. Following Daymont and Andrisani (1984) and Jann (2008), the gap in outcomes between groups may be expressed as a three way decomposition

$$(3) \quad (X_1 - X_2)\beta_2 + (\beta_1 - \beta_2)X_2 + (\beta_1 - \beta_2)(X_1 - X_2)$$

where the calculation is performed at mean values of the explanatory variables. The first term shows the endowment effect: the part of the gap that can be explained by differences in characteristics between groups. The second term shows the coefficients effect: the part of the gap that can be explained by differences in the returns to endowments between groups. The third term shows the interaction effect: the part of the gap that can be explained by observations in group 1 having a higher return on those characteristics for which they have higher means. Where an explanatory variable is a dummy, we normalise the effects relative to the grand mean to make the decomposition invariant to the choice of omitted category (see Yun 2005, Jann 2008). In the analysis section that follows, we provide three way decompositions in order to evaluate the source of the gap in the propensity to award good degrees, first between time periods, and secondly between different types of universities.

We make use of the three way decomposition instead of the more common two way decomposition because the two way decomposition is most often used in the discrimination literature, where the gap between groups is decomposed into an explained component and an unexplained component, the latter of which is usually attributed to discrimination. Whilst not uninformative, we believe that the three way decomposition we perform in this paper provides clearer insight into the determinants of differences between the different groups analysed.

### III. RESULTS

#### *Main results*

Estimated coefficients for the model are reported in Table 2. The simplest specification includes as explanatory variables only the average entry grades of students (three years before their cohort graduates) and the time dummies. The coefficient on entry grades is

positive and highly significant. The magnitude of the coefficient is intuitively plausible. An average intake which has three A levels each of grade A has a tariff score that (at midpoint) is 18% higher than an intake which has three A levels each of grade B, so the coefficient of 0.29 implies that a university with the former intake profile will graduate around 5% more of its students with good degrees.

[Table 2 here]

Few of the time dummies are significant. From 2008 onwards, however, there is a rising trend in the coefficients, and these become significantly different from zero in 2011 and 2012, suggesting some grade inflation. Indeed, by 2012, the proportion of good degrees awarded has risen by almost half a percentage point, other things being equal. While we cannot infer causality, it is interesting to note that this has coincided with a period of intensifying competition in the higher education sector in England, with undergraduate tuition fees rising from around £1000 per year to £3000 from the 2006 entry cohort, and later up to £9000 from the 2012 entry cohort.

In the remaining columns of the table, we add further explanatory variables. In the second column, it is readily observed that library spend, research score and student satisfaction all affect the dependent variable significantly and in an intuitively plausible direction. The impact of entry grades is almost unchanged from that reported in the first column. The time dummies are significant, but in all cases carry negative coefficients, suggesting that, relative to 2005 (the omitted year), grade inflation has been negative. However the magnitude of the coefficients has been declining in absolute terms since 2009, reflecting the pattern observed in column 1 of the table.

In column (3) of Table 2, we add a set of interaction terms between entry grades and the time dummies. Our intention here is to test the hypothesis that entry grades have become more important over time as a determinant of degree outcomes. This might be the case if sorting of students into degree programmes improves over time, strengthening peer effects and hence reinforcing the positive impact of higher education in institutions that recruit high quality intakes. The increased competitiveness of higher education in England, noted earlier, may well have contributed to prospective students becoming more selective in their choices of where and what to study.<sup>6</sup>

The coefficients on the interaction terms become significantly positive for interactions involving dummies for years from 2008 onwards, suggesting that there is indeed evidence for improved sorting. At the same time the absolute magnitude of the coefficients on the time dummies increases, though their significance falls, and there is no longer any significant difference between the coefficient for 2012 and zero. There is still some increase in the coefficients over the last part of the period under analysis, though this is confined to the 2010 through 2012 period. Overall, the evidence for grade inflation is therefore rather limited.

Finally, the last column of Table 2 shows the results of the model with all explanatory variables but no time dummies. These results are shown since the Blinder-Oaxaca decomposition between time periods which we perform next does not allow for the inclusion of time dummies. The results excluding the time dummies are broadly comparable to those in column (2) which include them.

Figure 2 plots the distribution of efficiency scores over time, omitting the outliers to highlight the main features. Median efficiency is high in all years (approximately 0.95), and shows no obvious trend. There may be some evidence that the distribution is becoming more compressed over time, which may indicate that increasing competition is reducing the differences in efficiency across institutions.

[Figure 2 here]

### ***Blinder-Oaxaca decomposition***

In Table 3, we report the results of a variety of three way decompositions. In the first column, we focus on the gap between the degree results variable in the years from 2005 to 2008 and in the years from 2009 to 2012<sup>7</sup>. The gap amounts to 5% in favour of the later period. Difference in endowments (especially research and entry grades) more than explains the total gap, while the later period in fact has worse returns to endowments than the earlier period, although this is due to a much lower constant term in the later period. The positive interaction between endowments and coefficients indicates that where the later period has a higher return, it is in variables in which the later period is better-endowed, and is significant for library spend and student satisfaction.

[Table 3 here]

In the remaining columns of Table 3 we provide decompositions for various specifications of the model, focusing this time on the gap between pre-1992 universities and other institutions.<sup>8</sup> Column (2) reports results without time dummies, column (3) adds the time dummies, and column (4) includes both time dummies and time-entry grade interactions. In 1992 a large number of higher education institutions were given university status for the first time. These were mainly polytechnics whose mission was focused on teaching, and they contrast sharply with institutions that had university status before 1992 in that the latter have missions that emphasise both teaching and research. While the distinction between pre-1992 and post-1992 institutions has blurred at the edges in the decades since this change, the two mission groups serving research intensive institutions (the Russell Group and the 1994 Group) both comprise only pre-1992 universities; the distinction between pre- and post-1992 institutions therefore remains relevant.

Before moving on to the discussion of the decomposition of the difference in outcomes between pre-92 and post-92 universities, we briefly refer back to Table 1. This reports descriptive statistics of the two time periods under consideration (2005 to 2008, and 2009 to 2012), and for the two groups of universities (pre-92 and post-92 universities). It can be seen that all values are higher for the later period than for the earlier period apart from the student-staff ratio which was unchanged. Similarly, pre-92 universities outperform post-92 universities in all the measures, although the differences are larger in some variables than in others. Figure 3 shows the percentage of good degrees over time by pre-92 and post-92 universities, showing that in both cases there is a rising percentage of good degrees over time. Figure 4 shows the efficiency scores of the two groups of universities; pre-92 universities are on average slightly more efficient than post-92 universities, and exhibit smaller variance both across institutions and over time.

[Figure 3 here]

[Figure 4 here]

In column (2) of Table 3 the gap in the proportion of good degrees awarded between the two groups of institutions is pronounced, at nearly 25% in favour of pre-92 universities. This is

roughly equally divided between the endowments, coefficients, and interaction effects. Among the university characteristics, it is differences in research quality and student satisfaction which have the most significant effects on the difference in good degrees; differences in entry grades also have a large but imprecisely measured effect. However, the really big difference between pre-92 and post-92 universities is in the return to university characteristics. There is a much higher return to entry scores in pre-92 universities, but a much lower return to student satisfaction. This is also reflected in the interaction terms: pre-92 universities have higher returns to characteristics in which they are well-endowed (such as entry grades). This suggests that students with good A level scores should attend pre-1992 universities because their talent is rewarded more highly there. This observation is clearly related to the sorting mechanism alluded to earlier.

In columns (3) and (4) of Table 3, the inclusion of the time dummies and interaction terms results in the interaction between endowments and coefficients playing the predominant role in explaining the gap between pre- and post-92 universities. Differences in endowments play almost no role in these specifications. We once again see a large effect of differences in the coefficient on entry scores, as well as on research scores, with pre-92 universities benefitting from higher payoff to both these characteristics than post-92 universities. This comes through in the interaction terms as well: for the characteristics that pre-92 universities excel in (entry grades and research scores), the return to the pre-92 universities is higher than to the post-92 universities. The conclusion is once again that good students do better in pre-92 universities than in post-92 universities.

### ***Additional results***

We next extend the basic model to analyse the conflicting findings of Achen and Thurnau (2009) and Popov and Bernhardt (2013). To recap, Achen and Thurnau find that a stronger market position would lead to less pressure for grade inflation, whereas Popov and Bernhardt argue that it is the better universities that would lower standards to give their students more favourable job market outcomes. We make use of two alternative measures of university quality: an indicator for whether the university is a pre-92 or a post-92 university, and an indicator for whether the university is a member of the Russell Group<sup>9</sup> of leading universities. These indicators are interacted with the year dummies, and the coefficients of

the interaction terms will shed light on whether it is higher-quality institutions that suffer more grade inflation over time, or vice versa.

[Table 4 here]

Table 4 reports these results; column (1) for the post-92 indicator, and column (2) for the Russell Group indicator; including both indicators in the same regression yields similar results. Controlling for all the other determinants of degree outcomes, post-92 universities award fewer good degrees than pre-92 universities. The magnitude of this coefficient is large; a post-92 university would typically award 14.6 percent fewer good degrees than a pre-92 university with identical characteristics. This is an important indicator that post-92 universities award fewer good degrees than pre-92 universities beyond the fact that they have fewer inputs into the production of good degrees (see Table 1). There is some suggestion that post-92 universities are more susceptible to grade inflation than pre-92 universities since the coefficients on the interaction terms are increasing over time; however, they are never statistically significant.

Column (2) of Table 4 reports the same exercise for Russell Group universities. The Russell Group indicator is positive and significant; Russell Group universities award more good degrees than other universities, even after controlling for other inputs into the production of good degrees. A member of the Russell Group would typically award 8.3 percent more good degrees than a non-member with the same characteristics. However, once again the interaction between the Russell Group indicator and the year dummies are not statistically significant, and in this case do not show any clear trend.

#### **IV. CONCLUSIONS**

Grade inflation grabs headlines. It provides empirical illustration of Goodhart's law – where a measure is used as a performance indicator, it becomes subject to manipulation (Goodhart, 1981). It is also a highly emotive subject – educated people are loath to see their qualifications devalued by subsequent relaxation of standards. However, ascertaining the extent to which grade inflation occurs is far from straightforward.

In this paper we have investigated the issue of whether grade inflation has occurred in higher education in the UK in recent years, using a stochastic frontier approach to take into account different levels of efficiency across universities. The evidence to support the existence of grade inflation is, at best, patchy. The quality of the student intake to universities has typically been rising over this period, and there have been changes in other factors that might reasonably be supposed to affect degree performance too. Moreover, the efficiency with which higher education institutions convert their inputs into outputs has remained fairly constant over time. The intensification of competition in the higher education sector appears to have resulted in an increase in the efficiency of the matching process between students and universities, and there is evidence that this too may have led to an improvement in outcomes. However, whilst better universities appear to be more lenient at grading, we find little evidence that different groups of universities exhibit different degrees of grade inflation over time.

The emotive nature of the content suggests that ours is unlikely to be the last word on the subject. This is good, not least because changes in the way university admissions are organised mean that there will in future be the opportunity for more rigorous consideration of some of the issues. In particular, as noted earlier, the introduction of adjustment offers the possibility of a discontinuity that could, with a few more years of data, be instructive.



---

## NOTES

<sup>1</sup> The Sunday Times (22 September 2013) refers to the scheme as an attempt to introduce more refined grade point averages (GPAs) alongside the traditional British degree classification – though it also reports that grade inflation is a concern in countries that use GPAs.

<sup>2</sup> Johnson (2003), Pressman (2007) and Babcock (2010) all suggest that a further real cost of grade inflation is that students tend to work less hard, and learn less, when good grades are relatively easy to come by.

<sup>3</sup> There are exceptions. Medicine and architecture, for example, are longer cycle degrees. Some students also undertake industrial placements or spend time studying abroad, both of which have the effect of lengthening their period of study. But, reflecting the fact that there is already a considerable measure of subject specialisation at upper secondary school in England and Wales, three years study for the bachelors degree remains the norm.

<sup>4</sup> Given the level of aggregation in our data we are unable to identify the reason for grade inflation, whether it is more lenient marking, technological change that allows different forms of assessment that enable higher marks, or some other reason. Identifying the reason for grade inflation is clearly an important issue for future research.

<sup>5</sup> The specifications used in our main table of results (Table 2) are re-run using conventional random effects methods to compare the results. These results are reported in the Appendix, and show that the two different estimation methods yield similar results.

<sup>6</sup> Future research may take note of a change in university admissions procedures in 2009 that provides a neat discontinuity and which may have affected the quality of matching between students and institutions. In that year, the Universities and Colleges Admissions Service, which coordinates applications to undergraduate study at British higher education institutions, introduced a system known as ‘adjustment’, which allows students to switch their choice of university *after* they get their A level results; this option is likely to be pursued by students who perform better than expected, and may lead to an improvement in sorting that was not possible in earlier years. Adjustment may have influenced the results observed in the present paper – the interaction term becomes significant only in the years after its introduction – but it has taken several years for this option to become widely appreciated and exercised. As recently as 2011, only a few hundred students exercised the option, but the numbers have risen rapidly since.

<sup>7</sup> Recall that this is measured on the frontier, so that we are studying a decomposition of  $y_1 - y_2 - (u_1 - u_2)$ .

<sup>8</sup> Note that the analysis of the gap between time periods in column 1 of Table 3 is, by necessity, conducted using a specification of the frontier model without time dummies. For the analysis of the gap between different types of university, the time dummies can be included; hence results from three different specifications of the model are reported in the remaining columns.

<sup>9</sup> “The Russell Group represents 24 leading UK universities which are committed to maintaining the very best research, an outstanding teaching and learning experience and unrivalled links with business and the public sector” (The Russell Group, 2013).

## REFERENCES

- ACHEN, A. and THURNAU, A. F. (2009). What are grades made of? *Journal of Economic Perspectives*, 23(3), 77-92.
- AIGNER, D.J., KNOX LOVELL, C.A. and SCHMIDT, P. (1977). Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics*, 6(1), 21-37.
- ANGLIN, P.M. and MENG, R. (2000). Evidence on grades and grade inflation at Ontario's universities. *Canadian Public Policy*, 26(3), 361-368.
- BABCOCK, P. (2010). Real costs of nominal grade inflation? New evidence from student course evaluations. *Economic Inquiry*, 48(4), 983-996.
- BELOTTI, F., DAIDONE, S., BARDI, G. and ATELLA, V. (2013). Stochastic frontier analysis using Stata. *The Stata Journal* 13(4), 719-758.
- BLINDER, A.S. (1973). Wage discrimination: Reduced form and structural estimates. *Journal of Human Resources*, 8(4), 436-455.
- BRACEY, G.W. (1994). Grade inflation. *Phi Delta Kappan*, 76(4), 328-331.
- CHAN, W., HAO, L. and SUEN, W. (2007). A signalling theory of grade inflation. *International Economic Review*, 48(3), 1065-1090.
- COMPTON, D.M. and METHENY, B. (2000). An assessment of grade inflation in higher education. *Perceptual and Motor Skills*, 90(2), 527-536.
- CORREA, H. (2001). A game theoretic analysis of faculty competition and academic standards. *Higher Education Policy*, 14(2), 175-182.

DAYMONT, T.N. and ANDRISANI, P.J. (1984). Job preferences, college major and the gender gap in earnings. *Journal of Human Resources*, 19(3), 408-428.

FREEMAN, D.G. (1999). Grade divergence as a market outcome. *Journal of Economic Education*, 30(4), 344-351.

GOODHART, C.A.E. (1981). Problems of monetary management: the UK experience. In A.S. Courakis (ed.) *Inflation, depression and economic policy in the West*, Washington: Rowman and Littlefield, 111-146.

GREENE, W. (2005). Reconsidering heterogeneity in panel data estimators of the stochastic frontier model. *Journal of Econometrics*, 126(2), 269-303.

HIGHER EDUCATION ACADEMY (2012). *A marked improvement: transforming assessment in higher education*, York: Higher Education Academy.

IZADI, H., JOHNES, G., OSKROCHI, R. and CROUCHLEY, R. (2002). Stochastic frontier estimation of a CES cost function: the case of higher education in Britain. *Economics of Education Review*, 21(1), 63-71,

JANN, B. (2008). The Blinder-Oaxaca decomposition for linear regression models. *The Stata Journal*, 8(4), 453-479.

JEWELL, R.T., MCPHERSON, M.A. and TIESLAU, M.A. (2013). Whose fault is it? Assigning blame for grade inflation in higher education. *Applied Economics*, 45(9), 1185-1200.

JOHNES, G. (2004). Standards and grade inflation. In G. Johnes and J. Johnes (eds) *International Handbook on the Economics of Education*, Cheltenham: Edward Elgar, 462-483.

JOHNES, G. and JOHNES, J. (2009). Higher education institutions' costs and efficiency: taking the decomposition a further step. *Economics of Education Review*, 28(1), 107-113.

JOHNES, G. and MCNABB, R. (2002). Academic standards in UK universities: more for less or less for more?. Paper presented at the Royal Economic Society Conference, 27 March.

JOHNSON, V.E. (2003). *Grade inflation: a crisis in college education*, New York, Springer-Verlag.

JONDROW, J., KNOX LOVELL, C.A., MATEROU, I.S. and SCHMIDT, P. (1982). On the estimation of technical efficiency in the stochastic frontier production function model. *Journal of Econometrics*, 19(2-3), 233-238.

KOLEVZON, M.S. (1981). Grade inflation in higher education: a comparative study. *Research in Higher Education*, 15(3), 195-212.

MARKS, D. (2002). Academic standards as public goods and varieties of free-rider behaviour. *Education Economics*, 10(2), 145-163.

MCKENZIE, R.B. and Tullock, G. (1981). *The new world of economics*, Homewood: Irwin.

OAXACA, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, 14(3), 693-709.

POPOV, S.V. and BERNHARDT, D. (2013). University competition, grading standards and grade inflation. *Economic Inquiry*, 51(3), 1764-1778.

PRESSMAN, S. (2007). The economics of grade inflation. *Challenge*, 50(5), 93-102.

ROJSTACZER, S. and HEALY, C. (2010). Grading in American colleges and universities. *Teachers College Record*, ID15928, accessed 23 September 2013.

SABOT, R. and WAKEMAN-LINN, J. (1991). Grade inflation and course choice. *Journal of Economic Perspectives*, 5(1), 159-170.

SMITH, J. and NAYLOR, R. (2001). Determinants of degree performance in UK universities: a statistical analysis of the 1993 student cohort. *Oxford Bulletin of Economics and Statistics*, 63(1), 29-60.

THE RUSSELL GROUP (2013), The Russell Group [online]. Available at <http://www.russellgroup.ac.uk/> [Accessed 9 January 2014].

TYMMS, P. and FITZ-GIBBON, C. (2001). Standards, achievement and educational performance: a cause for celebration?. In R. Phillips and J. Furlong (eds.) *Education, reform and the state: 25 years of politics, policy and practice*, London: Routledge Falmer, 156-173.

UNIVERSITIES UK (2007). *Beyond the honours degree classification: Burgess Group final report*. London: Universities UK.

UNIVERSITIES UK (2012). *Bringing it all together: Introducing the HEAR*, London: Universities UK.

WOSSMANN, L. (2003). Schooling resources, educational institutions and student performance: the international evidence. *Oxford Bulletin of Economics and Statistics*, 65(2), 117-170.

YUN, M.-S. (2005). A simple solution to the identification problem in detailed wage decompositions. *Economic Inquiry*, 43(4), 766-772.

TABLE 1  
DESCRIPTIVE STATISTICS.

Variable	Full sample (N = 826)		2005-2008 sample (N = 381)		2009-2012 sample (N = 445)		Pre-92 universities (N = 436)		Post-92 universities (N = 390)	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
% good degrees	61.190	10.839	59.268	10.641	62.835	10.748	68.239	9.040	53.309	6.312
Entry grades	296.165	77.583	281.981	73.528	308.310	78.973	350.733	64.499	235.162	32.328
Student-staff ratio	17.454	3.364	17.541	3.454	17.379	3.287	15.522	2.853	19.614	2.456
Library spend	767.440	321.398	672.411	277.500	848.801	334.135	894.844	364.778	625.007	178.167
Facilities spend	275.644	115.324	237.465	96.109	308.333	120.354	306.505	109.263	241.144	112.229
Research score	49.760	20.984	41.002	26.463	57.258	9.832	63.969	10.898	33.875	17.938
Student satisfaction	76.277	6.386	72.225	4.007	79.747	5.989	77.709	6.776	74.677	5.502

Notes: Entry grades are based on the UCAS points system in which an English pre-university qualification with three or four subjects is awarded 120 points for an A, 100 points for a B, 80 points for a C, 60 points for a D, and 40 points for an E. Library and facilities spend have been converted into real terms using the Consumer Price Index with base year in 2005. The research score is obtained from the Research Assessment Exercise and expressed as a percentage, while Student satisfaction is obtained from the National Student Survey and is the percentage of students who rated their course as satisfactory.

TABLE 2  
RESULTS OF THE TRUE RANDOM EFFECTS MODEL.

Dep var: ln(% good degree)	(1)	(2)	(3)	(4)
Ln(Entry grades)	0.290 (0.067)***	0.271 (0.034)***	0.167 (0.056)***	0.231 (0.075)***
Ln(Student-staff ratio)		-0.057 (0.036)	-0.054 (0.034)	-0.055 (0.032)*
Ln(library spend)		0.033 (0.017)*	0.026 (0.018)	0.002 (0.017)
Ln(facilities spend)		0.003 (0.017)	0.001 (0.016)	-0.001 (0.014)
Ln(research score)		0.028 (0.008)***	0.053 (0.018)***	0.030 (0.008)***
Ln(student satisfaction)		0.152 (0.079)*	0.159 (0.074)**	0.204 (0.042)***
2006	-0.011 (0.007)	-0.016 (0.008)**	-0.116 (0.173)	
2007	-0.019 (0.009)**	-0.039 (0.010)***	-0.326 (0.190)*	
2008	-0.021 (0.008)**	-0.042 (0.010)***	-0.350 (0.186)*	
2009	-0.006 (0.010)	-0.044 (0.011)***	-0.815 (0.292)***	
2010	-0.001 (0.011)	-0.041 (0.011)***	-0.878 (0.310)***	
2011	0.018 (0.011)*	-0.036 (0.017)**	-0.712 (0.303)**	
2012	0.043 (0.012)***	-0.016 (0.018)	-0.397 (0.307)	
Entry grades*2006			0.018 (0.030)	
Entry grades*2007			0.052 (0.033)	
Entry grades*2008			0.057 (0.032)*	
Entry grades*2009			0.135 (0.051)***	
Entry grades*2010			0.147 (0.054)***	
Entry grades*2011			0.119 (0.052)**	
Entry grades*2012			0.067 (0.053)	
<i>N</i>	826	826	826	826
Universities	115	115	115	115
$\sigma_u$	0.0826 (0.009)***	0.0768 (0.010)***	0.0774 (0.008)***	0.0695 (0.013)***
$\sigma_v$	0.0225 (0.008)**	0.0270 (0.007)***	0.0249 (0.006)***	0.0335 (0.009)***
$\lambda$	3.673 (0.017)***	2.846 (0.017)***	3.110 (0.013)***	2.0768 (0.021)***

Notes: \* significant at the 10% level; \*\* significant at the 5% level; \*\*\* significant at the 1% level. Estimation method is Greene's (2005) True Random Effects model with heteroskedastic-robust standard errors in parentheses.  $\lambda = \sigma_u/\sigma_v$ .

TABLE 3  
THREE WAY BLINDER-OAXACA DECOMPOSITION OF THE RESULTS.

		(1)	(2)	(3)	(4)
		Group 1 = after 2008	Group 1 = Pre-92	Group 1 = Pre-92	Group 1 = Pre-92
overall	Group 1	4.174 (0.019)***	4.277 (0.006)***	4.269 (0.013)***	4.275 (0.005)***
	Group 2	4.124 (0.017)***	4.027 (0.024)***	4.050 (0.040)***	4.044 (0.040)***
	difference	0.050 (0.025)**	0.249 (0.024)***	0.219 (0.043)***	0.230 (0.040)***
	endowments	0.071 (0.029)**	0.071 (0.021)***	0.008 (0.061)	0.012 (0.052)
	coefficients	-0.130 (0.047)***	0.106 (0.034)***	0.018 (0.044)	-0.022 (0.040)
	interaction	0.109 (0.055)**	0.073 (0.031)**	0.193 (0.066)***	0.240 (0.052)***
	endowments	Ln(Entry grades)	0.019 (0.008)**	0.026 (0.021)	0.009 (0.032)
	Ln(Student-staff ratio)	0.001 (0.001)	0.007 (0.012)	0.015 (0.011)	0.014 (0.014)
	Ln(library spend)	0.002 (0.006)	0.003 (0.008)	0.007 (0.009)	0.006 (0.009)
	Ln(facilities spend)	-0.001 (0.011)	-0.003 (0.006)	-0.002 (0.006)	-0.002 (0.006)
	Ln(research score)	0.057 (0.016)***	0.024 (0.007)***	-0.021 (0.030)	-0.014 (0.029)
	Ln(student satisfaction)	-0.007 (0.010)	0.014 (0.003)***	0.001 (0.005)	0.002 (0.005)
	time			0.000 (0.004)	0.004 (0.034)
	interact				-0.002 (0.036)
coefficients	Ln(Entry grades)	0.278 (0.632)	0.853 (0.450)*	1.680 (0.548)***	1.787 (0.401)***
	Ln(Student-staff ratio)	0.173 (0.150)	-0.022 (0.184)	0.130 (0.197)	0.121 (0.167)
	Ln(library spend)	0.423 (0.215)**	-0.105 (0.179)	0.004 (0.213)	-0.020 (0.173)
	Ln(facilities spend)	-0.012 (0.249)	0.112 (0.146)	0.151 (0.137)	0.146 (0.111)
	Ln(research score)	0.340 (0.319)	0.068 (0.051)	0.288 (0.137)**	0.467 (0.119)***
	Ln(student satisfaction)	1.350 (0.476)***	-0.802 (0.346)**	0.572 (0.666)	0.277 (0.579)
	Constant	-2.683 (0.823)***	-0.000 (0.537)	-2.801 (1.219)**	-2.791 (1.076)***
	time			-0.005 (0.004)	-0.001 (0.028)
	interact				-0.007 (0.028)



TABLE 3 (CONTINUED)  
THREE WAY BLINDER-OAXACA DECOMPOSITION OF THE RESULTS.

		(1)	(2)	(3)	(4)
		Group 1 = after 2008	Group 1 = Pre-92	Group 1 = Pre-92	Group 1 = Pre-92
interaction	Ln(Entry grades)	0.004 (0.010)	0.061 (0.032)*	0.121 (0.040)***	0.129 (0.029)***
	Ln(Student-staff ratio)	-0.000 (0.001)	0.002 (0.015)	-0.011 (0.016)	-0.010 (0.014)
	Ln(library spend)	0.016 (0.008)*	-0.005 (0.009)	0.000 (0.011)	-0.001 (0.009)
	Ln(facilities spend)	-0.001 (0.013)	0.006 (0.007)	0.007 (0.007)	0.007 (0.006)
	Ln(research score)	0.059 (0.055)	0.017 (0.012)	0.069 (0.033)**	0.113 (0.029)***
	Ln(student satisfaction)	0.031 (0.011)***	-0.007 (0.003)**	0.005 (0.006)	0.002 (0.005)
	time			0.000 (0.004)	-0.002 (0.024)
interact				0.001 (0.023)	
<i>N</i>		826	826	826	826
<i>N</i> _1		445	436	436	436
<i>N</i> _2		381	390	390	390

Notes: \* significant at the 10% level; \*\* significant at the 5% level; \*\*\* significant at the 1% level. Heteroskedastic-robust standard errors in parentheses. Each column of the table divides the sample into two groups. In column (1), Group 1 is from 2009 to 2012 and Group 2 is from 2005 to 2008. In columns (2) to (4), Group 1 is for pre-92 universities while Group 2 is for post-92 universities. The first panel shows the overall difference in the outcome variable (log percentage good degree) and decomposes it into endowment, coefficient and interaction effects. The rest of the table shows the detailed decomposition by variable. Time and interact are the subsumed effects of all time dummies and the interactions between time dummies and entry grades, where the time dummies and interaction terms are normalised relative to the grand mean (Jann 2008).

TABLE 4  
INTERACTING YEAR WITH MEASURES OF UNIVERSITY QUALITY

Dep var: ln(% good degree)	(1)		(2)	
	Coefficient	S.E.	Coefficient	S.E.
Ln(Entry grades)	0.134	(0.042)***	0.232	(0.055)***
Ln(Student-staff ratio)	-0.049	(0.028)*	-0.045	(0.035)
Ln(library spend)	0.025	(0.015)	0.025	(0.017)
Ln(facilities spend)	0.002	(0.012)	0.005	(0.014)
Ln(research score)	0.001	(0.022)	0.029	(0.010)***
Ln(student satisfaction)	0.109	(0.064)*	0.127	(0.120)
2006	-0.011	(0.008)	-0.016	(0.009)*
2007	-0.015	(0.010)	-0.035	(0.011)***
2008	-0.016	(0.010)	-0.035	(0.014)**
2009	-0.011	(0.012)	-0.037	(0.017)**
2010	-0.005	(0.012)	-0.033	(0.018)*
2011	-0.003	(0.017)	-0.026	(0.027)
2012	0.011	(0.019)	-0.004	(0.031)
Post-92 university	-0.146	(0.048)***		
Post-92 * 2006	-0.014	(0.015)		
Post-92 * 2007	-0.028	(0.018)		
Post-92 * 2008	-0.023	(0.016)		
Post-92 * 2009	0.001	(0.033)		
Post-92 * 2010	0.002	(0.033)		
Post-92 * 2011	0.018	(0.034)		
Post-92 * 2012	0.037	(0.036)		
Russell Group			0.083	(0.038)**
Russell Group * 2006			0.007	(0.012)
Russell Group * 2007			0.014	(0.019)
Russell Group * 2008			0.004	(0.014)
Russell Group * 2009			0.012	(0.016)
Russell Group * 2010			0.009	(0.015)
Russell Group * 2011			0.011	(0.018)
Russell Group * 2012			0.002	(0.018)
<i>N</i>	826		826	
Universities	115		115	
$\sigma_u$	0.0801	(0.010)***	0.076	(0.010)***
$\sigma_v$	0.0191	(0.009)**	0.026	(0.008)***
$\lambda$	4.181	(0.018)***	2.932	(0.017)***

Notes: \* significant at the 10% level; \*\* significant at the 5% level; \*\*\* significant at the 1% level. Estimation method is Greene's (2005) True Random Effects model with heteroskedastic-robust standard errors in parentheses.  $\lambda = \sigma_u/\sigma_v$ .

FIGURE 1. Box plot of percentage of good degrees (first class and upper second class honours) over time.




FIGURE 2. Box plot of efficiency scores over time (using the model of column 3 of Table 2).




FIGURE 3. Percentage of good degrees by pre-92 and post-92 universities.





FIGURE 4. Efficiency scores by pre-92 and post-92 universities.



APPENDIX  
RANDOM EFFECTS RESULTS FOR THE SAME SPECIFICATIONS AS TABLE 2.

Dep var: ln(% good degree)	(1)	(2)	(3)	(4)
Ln(Entry grades)	0.456 (0.031)***	0.337 (0.033)***	0.240 (0.062)***	0.324 (0.035)***
Ln(Student-staff ratio)		-0.074 (0.034)**	-0.073 (0.032)**	-0.082 (0.033)**
Ln(library spend)		0.039 (0.017)**	0.033 (0.019)*	0.004 (0.016)
Ln(facilities spend)		-0.001 (0.013)	-0.005 (0.013)	-0.005 (0.013)
Ln(research score)		0.032 (0.009)***	0.062 (0.020)***	0.031 (0.008)***
Ln(student satisfaction)		0.234 (0.098)**	0.219 (0.091)**	0.174 (0.041)***
2006	-0.008 (0.008)	-0.013 (0.009)	-0.141 (0.226)	
2007	-0.026 (0.010)***	-0.045 (0.010)***	-0.323 (0.211)	
2008	-0.027 (0.009)***	-0.047 (0.011)***	-0.292 (0.209)	
2009	-0.010 (0.010)	-0.053 (0.012)***	-0.870 (0.340)**	
2010	-0.009 (0.011)	-0.051 (0.013)***	-0.988 (0.362)***	
2011	0.008 (0.011)	-0.054 (0.020)***	-0.865 (0.361)**	
2012	0.025 (0.011)**	-0.041 (0.022)*	-0.592 (0.355)*	
Entry grades*2006			0.023 (0.040)	
Entry grades*2007			0.051 (0.037)	
Entry grades*2008			0.046 (0.037)	
Entry grades*2009			0.143 (0.059)**	
Entry grades*2010			0.164 (0.063)***	
Entry grades*2011			0.143 (0.062)**	
Entry grades*2012			0.097 (0.061)	
<i>N</i>	826	826	826	826
Universities	115	115	115	115

Notes: \* significant at the 10% level; \*\* significant at the 5% level; \*\*\* significant at the 1% level. Estimation method is random effects with heteroskedastic-robust standard errors in parentheses.